

This index encompasses all Apple II Technical Notes, File Type Notes, and Sample Code. Technical Notes are denoted by a four-letter category (e.g., IIGs for Apple IIGs Notes), File Type Notes by the letters FTN, and Sample Code by the letters SC.

\$C800 space	Misc 11	BADCTL error	SmPt 7
\$E100CB	IIGs 57	BADCTLPRM	SmPt 7
/RAM volume	GSOS 2, PDOS 15, 18, 21, 23, 26, 8	bank aligned memory	IIGs 78
3.5 drive differences	UDsk 5	bank-switched memory	see language card
3.5 ROM IIC	IIC 7	BASIC.SYSTEM	PDOS 24
5.25 drives	Pasc 17, PDOS 21	buffers	PDOS 9
40/80-column switch	IIC 7	battery RAM	FTN \$5A/0002
40 columns	IIC 2	BeginUpdate	IIGs 75
switching to	IIE 8	BELL (firmware routine)	IIE 10
40COL (soft switch)	IIE 8	BELL1 (firmware routine)	IIE 10
48K run-time system (Pascal)	IIGs 29	Binary II	FTN \$E0/8000, Misc 14
50 Hertz	Pasc 15	Binary Library Utility (BLU)	Misc 14
74LS244	Mous 2	BindInt	GSOS 9, IIGs 18
74LS245	IIE 2	bit-encoded slot configuration	IIGs 69
74LS251	IIE 2	BLOAD (BASIC.SYSTEM)	PDOS 24
80-column card	IIE 9	block device formatting	PDOS 5
80-column firmware	PDOS 11, 15	block size	UDsk 4
80-column GETCHAR routine	GSOS 2, IIGs 25	boot disk (minimal)	GSOS 1
80-column text page 2	Misc 10	boot disks, Pascal run-time	Pasc 10
80COL (soft switch)	IIE 3	booting	UDsk 2
80STORE (soft switch)	IIC 2, IIE 8, IIE 10	BootInit functions	IIGs 73
320 mode	IIE 3, IIGs 68, IIE 10	border color	IIGs 63
640 mode	IIGs 4	bottleneck procedures (QuickDraw II)	IIGs 34, 97
3200-color picture	IIGs 4	boundsRect	IIGs 80
65816 microprocessor	FTN \$C0/0002, \$C1/0002	BRK instruction	IIGs 1
0 prefix	IIGs 90	BRKVECTOR	IIGs 1
A/D input	ATLK 8, GSOS 10, SC 16	BSAVE (BASIC.SYSTEM)	PDOS 24
acceleration	IIGs 9	BUBIT location	PDOS 29
access privileges	IIE 10	buffer	
AccessPriv	ATLK 8, GSOS 4	QuickDraw II	IIGs 72
ActionNDA	SC 18	UnPackBytes	IIGs 94
ADB	SC 19	buffering	
address tables	IIGs 25, 26, 30, 91	keyboard	IIC 6
AddToQueue	IIGs 90	serial	IIC 6, IIGs 25
AddToRunQ	IIGs 24	buffers, BASIC.SYSTEM	PDOS 9
AIFF	IIGs 24	bug	
AIFF-C	FTN \$D8/0000	UniDisk 3.5	UDsk 3
AlertWindow	FTN \$D9/0001	ProDOS 8	PDOS 23
ALLOC_INT	IIGs 48, SC 5, 7	HyperCard IIGs	HCGS 2
ALLOC_INTERRUPT	GSOS 9	bus contention	IIGs 68, Misc 3
Allow Removal	IIGs 18	bus, SmartPort	SmPt 3
ALTMASK (softswitches)	SmPt 9	busy flag, system	IIGs 71
ALTZE (soft switch)	IIE 10	BusyBox	SC 2
ANI (soft switch)	IIGs 30, 68	cache	GSOS 7
ANIOFF (soft switch)	IIE 3	cache priority	GSOS 3
analog output impedance, sound	IIGs 29	caching	GSOS 3
animation	IIGs 9	windows	SC 7
AnimDemo	IIGs 70, SC 3	CalcMenuSize	IIGs 43
annunciators	SC 3	callbacks, HyperCard IIGS	HCGS 1
APPEND (BASIC.SYSTEM)	IIE 10	cards	IIGs 68
Apple /// emulation	PDOS 24	DMA	IIE 2
Apple 3.5 drive	Misc 2	physical dimensions	IIGs 28
Apple Desktop Bus	UDsk 5	CASSIN (softswitch)	IIE 10
Apple II High-Speed SCSI Card	see ADB	CASSO (softswitch)	IIE 10
Apple IIE	SmPt 5	CATALOG (BASIC.SYSTEM)	PDOS 24
Apple IIE Card (Macintosh LC)	IIC all	CD-ROM	PDOS 17
Apple IIE	IIE 10, Misc 2, Misc 7	recursive	SmPt 9
Apple IIE Workstation Card	IIE all	CD-RW disks	FTN \$B9, IIGs 71
Apple IIGs	ATLK 2, 4, 6, 7, PDOS 23	CDAs	FTN \$C7
Apple IIGs Programmers Workshop	IIGs all	CDev.Data file	FTN \$C7, IIGs 100, 86, SC 15
Apple Preferred Format	IIGs 20	CDevs	
Apple Sampled Instrument Format	IIGs 20	CHAIN (BASIC.SYSTEM)	PDOS 24
AppleDouble	FTN \$D0/0002, IIGs 27	ChangePath	GSOS 13
AppleShare activity arrows	FTN \$D5/0007	changing resolution	IIGs 4
AppleShare volumes	FTN \$E0/0002-3	channels, SCC	IIGs 18
	IIGs 5	character devices	GSOS 4
	FTN \$B6, \$C7	status	GSOS 13
	ATLK 8,	character sets	IIE 10
	PDOS 17, 21, 22, 23, SC 16, 18, 22	characters, waiting	GSOS 13
	FTN \$E0/0001	chExtra	IIGs 72
	Misc 11	clamping (mouse)	Mous 7
	Mous 6	ClampMouse	Mous 1
	Misc 9	Claris	FTN \$19, \$1A, \$1B, \$50/8010,
	ATLK all, IIE 10		Misc 4, 5
	FTN \$BB, IIGs 18, 26, 77	Classic Desk Accessories	see CDAs
	see ASP	CLD instruction	PDOS 12, 22
	FTN \$50/8010	ClearMouse	Mous 3
	Misc 4	ClearScreen	IIGs 72
	FTN \$19	CLI instruction	PDOS 12
	FTN \$1B	clicks, multiple	IIGs 84
	FTN \$1A, 5	clipping	IIGs 99
	GSOS 10	clipping buffer	IIGs 72
	FTN \$B3	clipping	IIGs 80
	SC 1	clipRgn	IIGs 80
	IIGs 30	clock	IIE 10
	FTN \$B0, \$B5, IIGs 20, 33	clock	PDOS 1, 11, IIE 10
	IIGs 6	close	IIGs 53
	FTN \$D5/0007, \$D8/0002	CloseAllNDAs	IIGs 4, 53
	ATLK 5	CloseNDABYWinPtr	IIGs 71
	IIGs 77	CloseResourceFile	IIGs 71
	PDOS 23	CloseView	IIGs 91
	ATLK 1, 2	ClrHeartBeat	IIGs 55
	FTN \$D8/0001	CLRVLINT (soft switch)	IIGs 49
	FTN \$D8/0000, \$D8/0001	CMReleaseResource	IIGs 60
	IIGs 58	code resources	IIGs 86
	SC 22	attributes	IIGs 81
	IIGs 17	color table scrap	IIGs 99
	ATLK 2, PDOS 26	color tables	IIGs 63, 97
	PDOS 19	window	IIGs 98
	IIE 3	color values (RGB)	IIGs 63
	SC 16	command keys, NDAs	IIGs 71
	IIGs 63	comment, resource type	IIGs 76
	IIGs 62	CompactMem	IIGs 71
	IIGs 16	configuration files	ATLK 8
	PDOS 29	CONTROL	SmPt 6
	FTN \$01	CONTROL ^	IIGs 65

Control key	IIGs 58	escape key	Misc 10
Control Manager	IIGs 26, 81	Event Manager	IIGs 24, 26, 71, 91
Control N Monitor command	IIGs 25	events, menu	IIGs 24
Control Panel CDA	IIGs 26	ExpandPath	GSOS 10
Control Panel control jumper	IIGs 30	ExpressLoad	IIGs 66
Control Panel NDA	FTN \$C7	extended controls	see controls, extended
control init message	IIGs 81	extended file	PDOS 25
control panel devices	see CDeva	extended serial port firmware	IIGs 50
control record	IIGs 81	external commands and functions	see XCMD
control template	IIGs 81	EZ Backup	FTN \$E0/8006
controls, custom	IIGs 81, 86, SC 4, 9	fakeModalDialog	SC 20, 22, 9
controls, extended	IIGs 81	fakeMouse	IIGs 85
controls, inactive	IIGs 24	fastPort (QuickDraw II)	IIGs 72
controls, tracking	IIGs 84	Fatal System Error, \$0512	IIGs 24
convert (resource attribute)	IIGs 86	FFGeneratorStatus	IIGs 37
ConvSeconds	FTN \$D8/0000	FFStartSound	IIGs 11, 37
cool math (quadratic)	Misc 12	FFStopSound	IIGs 37
CopyFixels	IIGs 72	File System Translators	see FSTs
COUT routine	Misc 6	File Type Descriptors	FTN \$42
COUTI routine	IIGs 25	file level	GSOS 13, IIGs 71
CPU cycle	IIE 2	file transfer	Misc 14
CREATE (ProDOS 8)	UDsk 3	file types	PDOS 19
CROW0 and CROW1 signals	IIGs 21	filenames	GSOS 8
CtlNewRes	IIGs 4	fileSysID	GSOS 11
current resource application	IIGs 83	FileListSessions	PDOS 21
cursor		FilterProc	IIGs 38
QuickDraw II	IIGs 39, 76, 85	FindControl	IIGs 84
resource format	IIGs 76	Finder	FTN \$42, \$CA
shielding	IIGs 24	fInWindowOnly (flag bit)	FTN \$C7
text	IIGs 65	Firmware Reference Updates	IIGs 25
cursor manipulation	IIGs 85	firmware	
custom controls, extended	IIGs 81	80-column	IIGs 25
Custom.Control	SC 4	ID bytes	IIGs 25
Custom.Window	SC 5	mouse	IIGs 25
CYA	IIGs 21	serial port	IIGs 16, 25, 26, 50
cycle timings, 65816	IIGs 70	FixAppleMenu	IIGs 24
Darts	SC 17	Fixed data type	IIGs 79
data base, AppleWorks	FTN \$19	FixMenuBar	IIGs 43
data bus	IIE 2	floating point numerics	SC 21, see SANE
dates, ProDOS	PDOS 28	Floyd	Misc 14
Devex archived volume	FTN \$E0/8004	Flush	GSOS 13
DCE devices	IIGs 30	folder selection	IIGs 96
DEALLOC_INTERRUPT	IIGs 18, PDOS 12	Font Manager	IIGs 15, 26, 46, 53, SC 14
debugger (custom)	IIGs 1	font families	IIGs 41, 67
deProc, window	IIGs 42	FontReport	SC 14
DeleteFromQueue	IIGs 24	fonts	FTN \$C8
DeleteMItem	IIGs 24	custom	ImWr 1
dereferencing handles	IIGs 90	large	IIGs 15, GSOS 1, IIGs 4, 41, 67, SC 14
Desk Manager	IIGs 26	FORMAT	SmPt 9, UDsk 4
desk accessories	FTN \$B8, \$B9, IIGs 53, 71, SC 19	Format (GS/OS call)	GSOS 11
and menus	HCGS 2	formatter	
DeskShutDown	IIGs 53	Pascal	Pasc 12
DESTROY (ProDOS 8)	FTN \$01, PDOS 23	ProDOS	PDOS 5
DEVCONT location	PDOS 20, 21	formatting	
development tool (8 bit)	SC 21	Macintosh disks	UDsk 4
DEVICE SELECT signal	SC 4	ProDOS disks	PDOS 16
device names	GSOS 4	FotoFile	FTN \$08/0-3FFF
device status, character	GSOS 13	FPI	IIGs 21
device types (SmartPort)	SmPt 4	Frac data type	IIGs 79
devices		fractions (cool diagrams)	IIGs 79
character	GSOS 4, 13	FrameBgn	IIGs 24
identifying	PDOS 20	free-form synthesizer	IIGs 37
ProDOS	PDOS 20, 21	FREEBUF routine (BASIC.SYSTEM)	PDOS 9
DEVLIST table	IIGs 68	FreeMem	IIGs 51
DEVSEL signal	IIGs 26, 38, 91, SC 6	FrontWindow	IIGs 91
Dialog Manager	ATLK 1, GSOS 12, IIGs 69	FSTs	FTN \$BD
DInfo	see DMA	function pointer table (FPT)	IIGs 101
direct memory access	GSOS 6, 7	FWEntry	IIGs 69
direct page (GS/OS)	GSOS 4	fZoomed (flag bit)	IIGs 24
directory structure	PDOS 21, 23	game I/O	IIE 10
Disk II	IIGs 30	GET_FILE_INFO (ProDOS 8)	UDsk 3
disk port soft switches	IIGs 25	GETBUF routine (BASIC.SYSTEM)	PDOS 9
disk sector format (3.5)		GetDirEntry	GSOS 4, 13
disks		GetFanNum	IIGs 41
3.5, Pascal	Pasc 16	GetFirstDItem	IIGs 38
3.5, single-sided	IIC 8	GetFrameColor	IIGs 45
Macintosh	UDsk 4	GetFSTInfo	ATLK 1
switching	UDsk 4	GetInfo	ATLK 1, 2, 4
display screens	IIE 10	GetInterruptState	IIGs 24
DisposeAll	IIGs 17	GetIntInfo	IIGs 25
DisposeHandle	SC 6	GetLevel	GSOS 13, IIGs 53
DLog	IIGs 30	GETLN (firmware routine)	IIE 10, IIGs 65
DMA bank register	IIE 2, IIGs 21, 68	GETLN buffer	PDOS 1
DMA	IIGs 21	GETLN1 (firmware routine)	IIE 10
DMA register (\$C037)	IIGs 11, 9	GETLN2 (firmware routine)	IIE 10
DOC chip	IIGs 19	GetLoadSegInfo	IIGs 66
DOC mode	IIGs 23, 53	GetMenuTitle	IIGs 60
DOC RAM (sound)	PDOS 3	GetMItem	IIGs 24, 60
DOS 3.3	PDOS 2	GetMItemName	IIGs 24
DOSCMD (BASIC.SYSTEM)	IIE 3	GetModeBits	IIGs 18
double hi-res page 2	FTN \$08/0-3FFF	GetNewID	IIGs 71
double high-resolution graphics	IIC 2, IIE 3, 10	GetNextEvent	IIGs 71, 91
packed	FTN \$08/4001, IIC 2, IIE 3	GetPenState	IIGs 44
DOWNLOAD	UDsk 4	GetPortRect	IIGs 80
DragWindow	IIGs 71	GetPrivileges	SC 18
DrawMenuBar	IIGs 60	GetUserPath	ATLK 8, GSOS 10
DrawPicture	IIGs 46, 72	GetVector	IIGs 1
DrawName	GSOS 12	GetWaitStatus	GSOS 13
drive, Apple 3.5	see Apple 3.5 drive	GetWAP	IIGs 73
drivers	FTN \$BB	GetWTitle	IIGs 61
GS/OS	IIGs 69, 100, GSOS 6	Golden NDA Guideline (ask Matt)	IIGs 53
port	IIGs 36	GPB (General-Purpose interface bus)	IIE 7
printer	IIGs 93, see printer drivers	GR (soft switch)	IIE 3, IIGs 29
ITS devices	IIGs 30	GraOff	FTN \$B3
ITS Tools and Libraries	SC 20	GraOff	FTN \$B5
Dvorak keyboard	IIC 4	grafPort	IIGs 35
DYN_SLOT_ARBITER	GSOS 4, IIGs 69	current	IIGs 24, 82
dynamic segments	IIGs 22, 56	printing	IIGs 35, 80, 91, 93
Dynamo	SC 21	grafProc	IIGs 34
Easy Access	IIGs 91	Graphic Disk Labeler document	FTN \$53/8002
EH?	IIGs 83	graphics	IIGs 70
EJECT	SmPt 2	ground noise	IIGs 68
Elms (SANE)	IIGs 8	GS/OS application	FTN \$B3
EMStatus	IIGs 71	GS/OS	
emulation mode (65816)	IIGs 90	versions	GSOS all, PDOS 23, 27, SC 16, 18, 22
EndFrameDrawing	SC 5	GS/OS aware	IIGs 49, 69, 93, 100, FTN \$BD,
EndInfoDrawing	IIGs 3	GS/OS drivers	FTN \$B3, \$B5
EndOPicture	IIGs 46	GS/OS shell application	FTN \$BB, IIGs 100
EndUpdate	IIGs 75	half-dot shift	FTN \$B5
enhanced IIE ROMs	IIE 8	HALT (SANE)	IIE 3
ENVB, (soft switch)	IIE 9	handles, dereferencing	IIGs 7
EraseDisk	IIGs 11	hardware access	IIGs 90
error codes, QD Aux	IIGs 24	hardware reference updates	IIE 10
ERRORDEATH macro	IIGs 33	hardware, serial port	IIGs 30
errors		heartbeat tasks	IIGs 30
Pascal booting	Pasc 10	HFS (Hierarchical File System)	FTN \$B6, SC 3
ProDOS devices	PDOS 21	high-resolution graphics	UDsk 4
SmartPort	IIGs 25		FTN \$08/0-3FFF

packed	FTN \$08/4000, IIGs 29	maxWidth	IIGs 72
HiliteMenu	IIGs 60	Medley document	FTN \$54/DD3E
HIRES (soft switch)	IIE 3, IIGs 29	Mega II	IIGs 32, 68
HPIB (Hewlett-Packard interface bus)	IIE 7	Mega II video counters	IIGs 39
hybrid applications	PDOS 27	Memory Expansion Card	IIE 10, MemX 1
HyperCard IIGS		Memory Manager	IIGs 17, 26
bugs	HCGS 2	Memory Peeker desk accessory	IIGs 25
HyperCard IIGS Script Language Guide	HCGS 1	memory compaction	IIGs 57
HyperStudio sound	FTN \$D8/8001	memory expansion	IIC 5, IIGs 21
I/O redirection (BASIC.SYSTEM)	PDOS 4	memory expansion slot	IIGs 21
I/O SELECT signal	IIE 4	memory ID	see User ID
I/O STROBE signal	IIE 4	memory management	GSOS 3
I/O subroutines	IIGs 2	and interrupts	IIGs 4, 17, 22, 51, 52, 57, 78, 91,
icon scrap	IIGs 99	memory purging	PDOS 26, 27
icons	FTN \$CA	Menu Manager	IIGs 57
ID bytes	Misc 2	menu bars, system vs. window	IIGs 26, 43, 60, 81, SC 12
mouse	Mous 5	menu events	IIGs 60
UniDisk	UDsk 2, 7	menu strings	IIGs 24
ID nibble	PDOS 21	menu titles	IIGs 5
identification of CPUs	Misc 2	menu, custom	SC 12
identifying devices	PDOS 21	MenuGlobal	IIGs 60
idv values	IIGs 35, 93	MenuNewRes	IIGs 4
idle events	IIGs 84	Menus	SC 12
IDROUTINE routine	Misc 7	MenuSelect	IIGs 60
IDSEARCH routine	Pasc 14	MenuStartUp	IIGs 60
IEEE-488 card	IIE 7	MessageByName	IIGs 89
IFF	FTN \$D8/0001	MessageCenter	FTN \$B3, \$B5, IIGs 89
IIC	IIC all	MFS (Macintosh File System)	UDsk 4
slots	IIE all	MIXED (soft switch)	FTN \$D7
versions	IIC 7	MIDI drivers	PDOS 27
IIC Plus	IIC 7, SmPt 7	MIDI Synth	IIGs 54
IIE	IIE all	MIDI Tools drivers	IIGs 54
IIGs	IIGs all	MIDI tools	FTN \$BB
ImageWriter II	ATLK 9	Miscellaneous Tools	IIGs 23, 54
ImageWriter	ImWr 1, SC 9	MIXED (soft switch)	IIGs 26, 94
ImageWriter LQ	ATLK 9	MLINACTV location	IIE 3
information bar	IIGs 24, 3	NMStartUp	ATLK 7, PDOS 23
INH line	IIE 5, IIGs 32	modal dialog	IIGs 17, 53, PDOS 27
INIT	SmPt 2, 9, UDsk 4	modal windows in NDAs	IIGs 91, SC 6
initialization files	FTN \$B6, \$B7	ModalDialog	IIGs 71
InitialLoad	IIGs 66, 73	mode switching (320/640)	IIGs 38
InitialLoad2	IIGs 86	mode, mouse	IIGs 4
InitMouse	Mous 2	modeless dialog	Mous 3
INSTALL_DRIVER	GSOS 6	modifiers	SC 6
Installer	IIGs 64	MoMan	IIGs 58
InstallFont	IIGs 15	Monitor	SC 22
InstallNDA	IIGs 71	monochrome hi-res graphics	IIGs 25, 26
InstallTimer	ATLK 5	Moofi (TM)	IIGs 29
instruction timing, 65816	IIGs 2	mouse	SC 10
instrument	FTN \$D6	mouse flags	IIGs 81
Integer Math	IIGs 26, 79, SC 13	motherboard	IIE 9
Integrated Woz Machine	see IWM	Mouse Keys	IIGs 91
INTEN (soft switch)	IIGs 49	mouse button status	IIGs 25
interchange formats	FTN \$D8/0000	mouse card	IIC 1, Mous 5
interfaces	IIE 7	mouse firmware	IIGs 25
interlace mode	Misc 17	mouse	IIE 10, mous all
interleave	UDsk 4	behavior	IIC 1
international country codes	IIGs 76	clamping	Mous 7
interrupt state record	IIGs 24	identification	Mous 5
interrupt status, SmartPort	SmPt 2	positioning	IIGs 85
interrupt, non-maskable	IIGs 1	scaling	IIC 1
interrupts	ATLK 2	mouse mode	Mous 3
SCC	IIGs 18, 71	MouseText	Mous 6
serial	IIC 6	mouseUp events	IIGs 71
unclaimed	IIE 8, IIGs 25, 70, GSOS 9,	MOVE_INFO	GSOS 6
	Mous 1, 4, PDOS 12	MPW IIGs	IIGs 33
IntSource	IIE 19	MSLOT location	IIGs 16, Misc 3
InvalRect	IIGs 24, 75	MTR (BASIC.SYSTEM)	PDOS 24
InvalRgn	IIGs 10, 75	multimedia drivers	FTN \$BB
Invisible bit	ATLK 6	Music Construction Set	FTN \$D5, \$D6
IOSEL signal	IIGs 68	music sequence	FTN \$D5
IOSTB signal	IIGs 68	NDAs	FTN \$B8, IIGs 71, SC 19
IOUDIS (soft switch)	IIGs 68	network volumes	ATLK 8
IRQ signal	IIGs 74	New Desk Accessories	see NDAs
item draw routine (List Manager)	IIGs 30, Misc 6, UDsk 5	NewControl2	IIGs 81
IWM	IIGs 91	NewHandle	IIGs 17
J (middle initial)	SC 7	NewMenu	IIGs 60
Jiffy.Windows	IIE 10	NewMenuBar2	IIGs 60
joystick connector	IIE 6	NEWVIDEO (soft switch)	IIGs 70
joystick	IIGs 30	NewWindow	IIGs 24, 47
jumper, Control Panel (S1)	IIC 6	NewWindow2	IIGs 24, 3, 82, SC 2, 9
keyboard buffering	IIE 10	NextMember2	IIGs 24
keyboard layout	IIGs 58	Nifty, Mr.	FTN \$B9
keyboard modifiers register	remove	NIL in resource maps	IIGs 83
keyboard, Dvorak	IIC 4	NMI	IIGs 1
keyboard, foreign	IIC 3	NMI signal	IIGs 68
keyboard, reading from CDAs	IIGs 71	no-special-memory attribute	IIGs 52
KEYMODREG (soft switch)	IIGs 58	no-wait mode	GSOS 13
Krunching	Pasc 10	non-maskable interrupt	IIGs 1
KVERSION location	PDOS 23	Note Sequencer	IIGs 23
language card	PDOS 12	Note Synthesizer	IIGs 19, 23
LaserWriter	ATLK 9, IIGs 41, 67, SC 9	notification procedures	FTN \$B6, GSOS 12, 7
LaserWriter font mapping	IIGs 67	NWSC color	IIE 3
LCENK2 (soft switch)	IIGs 30	NuFX	FTN \$E0/8002, Misc 14
Lee, Jane and 'rutabaga' incident	Misc 13	Object Module Format	FTN \$BC, IIGs 66
LEFromScrap	IIGs 59	offscreen ports	SC 20
LEIdle	IIGs 84	OMF	FTN \$BC, IIGs 66
LEKey	IIGs 84	OMF KIND field	IIGs 52, 78
level	see file level	ONLINE (ProDOS 8)	IIC 5, PDOS 8, 21, 23
LGetPathname2	IIGs 71	onlyGetSelection	IIGs 92
Line Edit controls	IIGs 81	OCM routine	IIGs 51
Line Edit	IIGs 26, 81, SC 8	Open	SC 16
List Manager	IIGs 24, 26, 38, 74, SC 17, 8	OpenPicture	IIGs 72
list controls	IIGs 38	OpenPoly	IIGs 72
list members, inactive	IIGs 24	OpenPort	IIGs 91
List.Line.Edit	SC 8	OpenRgn	IIGs 72
Lister	SC 9	option_list	FTN \$BD, GSOS 4
lists in dialogs	IIGs 38	OS_BOOT location	PDOS 27
load files	FTN \$BC	OSShutdown	GSOS 2, IIGs 49
loaders	IIGs 22, 52, 66	out-of-memory queue, corrupted	IIGs 24
LoadOneTool	IIGs 100, 53	out-of-memory routine	IIGs 51, 78
LoadResource	IIGs 83	P-Machine	Pasc 10
LoadSegName	IIGs 22	PackBytes format	IIGs 27
LoadTools	IIGs 100	PackBytes	FTN \$08/0-3FFF, \$08/4000, \$08/4001,
locInfo	IIGs 6, 80	packed super-hires image	\$C0/0001, \$C0/0002, IIGs 94
longStatText2 (dialog item type)	IIGs 91	packets, SmartPort	FTN \$C0/0001
lowercase	GSOS 8	paddles	SmPt 8
M2B0 signal	IIGs 68	PAGE2 (soft switch)	IIE 6
M2SEL signal	IIGs 68	PaintWorks packed picture	IIE 3, IIGs 68
MACHD location	PDOS 11, 15	PAL timing chip	FTN \$C0/0000
Macintosh Audio Compression and Expansion (MACE)	FTN \$D8/0001	PAP	IIE 2
Macintosh disks	UDsk 4	PAPOpen	ATLK 3
Manners, Missed	PDOS 26	PAPStatus	ATLK 9
mask scrap	IIGs 99	Parallel Interface Card	IIE 7, Misc 16
Math	SC 13	parameter blocks, controls	IIGs 81
MaxBlock	IIGs 51	parameter blocks, ProDOS	PDOS 22
		ParamText	IIGs 91
		Pascal 1.1 firmware protocol	IIGs 16, Misc 8

Pascal and MouseText	Mous 6	resource formats	IIGs 76
Pascal area (on ProDOS volume)	PDOS 25	resource search depth	IIGs 83
Pascal	Pasc all	resource search order	IIGs 83
48K run-time system	Pasc 15	resource search path, CDevs	FTN \$C7
Patch	Pasc 17	resources	SC 2
run-time boot disks	Pasc 10	as templates	IIGs 83
Pascal Profile manager	PDOS 25	ResourceStartUp	IIGs 53, 71
Pascal run-time system	Pasc 10	Restart	IIGs 52, 73
Pascal volumes	Pasc 16	ReturnStat (HyperCard IIGs parameter)	HCGS 1
Patch (Pascal)	Pasc 17	Revision B motherboard	IIE 3
patching, tool dispatcher	IIGs 87	ROM Disk	IIGs 25
patching, tools	IIGs 101	ROM revisions	IIGs 26
pathnames	GSOS 4	ROM, IIE enhanced	IIE 3
patterns	FTN \$C0/0002	rowBytes	IIGs 80
window backgrounds	IIGs 62	RS-232-C	IIGs 30
FDLTRIG (softswitch)	IIE 10	RS-422	IIGs 30
PEEK (AppleSoft function)	Misc 11	RSHMEM routine (BASIC.SYSTEM)	PDOS 9
pen pattern	IIGs 6	RST signal	IIGs 68
pen state record	IIGs 44	rStyleBlock	IIGs 99, FTN \$50/5445
permanent initialization files	FTN \$B6, \$B7	rTextForLETextBox2 (resource type)	IIGs 24
PFI	PDOS 21	RTI instruction	PDOS 12
PH0 signal	IIE 4, IIGs 68	run queue	IIGs 24, FTN \$B6
PH1 signal	IIE 4	run-time system, Pascal	Pasc 10
PH2 signal	IIE 4, IIGs 21, 68	rVersion	IIGs 76
PicComments	IIGs 97, SC 9	samled sound	FTN \$D8
PICt	FTN \$C1/0001, IIGs 27, 46	sampling frequency	IIGs 37
picture (QuickDraw II)	IIGs 72, FTN \$C1/0001	SANE	FTN \$1B, IIGs 26, 7, 8, PDOS 26, SC 13
Macintosh	IIGs 46	scan line control bytes	IIGs 97
picture comments	IIGs 97	scanline counter	IIGs 39
picture data format	IIGs 46	SCC (chip)	IIGs 18
Picture, QuickDraw II	IIGs 27	SCC ARGB soft switch	IIGs 49
PINIT entry	IIGs 16	SchAddTask	IIGs 71
pinouts	IIE 7	scheduler	IIGs 16, 26, 71
Pixel Map Tools	SC 13, 20	Scrap Manager	IIGs 26, 53
pixelMap2Rgn	SC 20	scrap types	IIGs 99
PlaySound	SC 10	screen dump	IIGs 27
PMCloseSession	ATLK 3	screen holes	Misc 1
PMLoadDriver	IIGs 77, SC 11	screen image	IIE 10
PMSetPrinter	ATLK 3	packed	FTN \$C1/0000
PMUnloadDriver	IIGs 77	scripts, Installer	FTN \$C0/0001, IIGs 27
polygons	IIGs 72	ScrollRect	IIGs 64
Pop-Up Menu controls	IIGs 81	SCSI	IIGs 72
port drivers	IIGs 36	SmPt 5, 9	SmPt 5, 9
ports, offscreen	SC 20	segments	IIGs 66
portSCB	IIGs 80	dynamic	IIGs 22, 86
PostMouse	IIGs 85, Mous 1, 3	selection (Text Edit)	IIGs 92
PostScript	IIGs 67, 97	selector/dispatcher (ProDOS 8)	see QUIT code
power consumption, cards	IIGs 68	self test	IIE 9, IIGs 95
power-up byte	MemX 1	SendQueue	IIGs 16
POWERUP location	IIGs 49	separators	GSOS 4
PPtPort	IIGs 80	SerFlag	IIGs 18
PREAD (firmware routine)	IIE 6, 10	serial buffering	IIE 6, IIGs 25
prefixes	GSOS 10	serial controller chip	IIGs 18
preload (resource attribute)	IIGs 86	serial interrupts	IIE 6
Prevent Removal	SmPt 9	serial port firmware	IIGs 16, 25, 26, 50
PrGetUserName	IIGs 77	serial port hardware	IIGs 30
Print Manager drivers	FTN \$BB	serial ports	IIE 10
Print Manager	IIGs 35, 77, 93, SC 11, 9	ServeMouse	Mous 1, 4
print loop	IIGs 93	SET_DISKSW	GSOS 12, 7
print record	IIGs 35, 93, SC 11	SET_DOWN_ADDR	Udek 4
Printer Access Protocol	see PAP	SET_FILE_INFO	PDOS 29
printer driver	IIGs 35, 93, 97, SC 11	SET_SPEED	GSOS 6
Printer_Setup file	IIGs 77	SetArcRot	IIGs 6
printing	IIGs 35	SetAutoKeyLimit	IIGs 24
background	IIGs 16	SetBufDims	IIGs 72
transparent	ATLK 4, 97	SetContentOrigin	IIGs 47
printing text	IIGs 93	SetCtlMoreFlags	IIGs 81
printX subrecord	IIGs 93	SetCtlTitle	IIGs 81
ProDOS 8	ATLK 2, PDOS 4, FTN \$01	SetCurResourceFile	IIGs 83
bugs	GSOS 2, 8, SC 21	SetDataSize	IIGs 47
ProDOS file system limitations	PDOS 23	SetDefaultTPT	IIGs 101
ProDOS Filing Interface	GSOS 4	SetDTR example	IIGs 50
ProDOS	see PFI	SetFrameColor	IIGs 98
PrRecord.Spy	PDOS all	SetGraProc	IIGs 34, 35
PrValidate	SC 11	SetInfoDraw	IIGs 3
ptrCheck	IIGs 93	SETINTC3ROM (soft switch)	GSOS 2
PTRIG (soft switch)	SC 9	SetInterleave	IIGs 25
ptrToPixImage	IIE 9	SetInterruptState	IIGs 24
ptrToPixImage	IIGs 80, 91	SetIntInfo	IIGs 25
purgeable memory	IIGs 91	SetLevel	IIGs 26, IIGs 53
purging memory	IIGs 51	SetMemlFlag	IIGs 60
PWRITE entry	IIGs 78	SetModeBits	IIGs 25
Q Monitor command	IIGs 16	SetMouse	IIE 1, Mous 3, 4
Q6 and Q7 (IWM)	IIGs 25	SetMTitleStart	IIGs 5
QDStartUp	IIGs 30	SetOrigin	IIGs 80
QuickDraw II Aux, error codes	IIGs 72	SetOutBuff	IIGs 16
QuickDraw II Auxiliary	IIGs 24	SetPenState	IIGs 44
QuickDraw II	IIGs 24, 46, 53, 75	SetPort	IIGs 24
tutorial	IIGs 4, 10, 26, 34, 44, 6, 72, 80, 91	SetPrivileges	SC 18
QuickDraw II picture	IIGs 80	SetResLoad	IIGs 83
QUIT code	FTN \$C1/0001	SetResourceFileDepth	IIGs 83
R/W line	PDOS 23, 23	SETSLC3ROM (soft switch)	GSOS 2
R/W* line	IIE 2	SetStdProc	IIGs 34, 35
RAM Disk	IIE 2	SetSysBar	IIGs 60
RAMRD (soft switch)	GSOS 2, PDOS 8	SetTSPtr	IIGs 73
RAMWRT (soft switch)	IIE 3, IIGs 68	SetVector	IIGs 1
xComment	IIGs 76	SetWAP	IIGs 73
RDALZP (soft switch)	IIGs 25	SetWTitle	IIGs 61
RDDHRES (softswitch)	IIE 10	SetZoomRect	IIGs 24
RDIODUIS (soft switch)	IIE 10, IIE 9	SFMultiGet2	SC 18
RDY line	IIE 4	SFPGetFile2	IIGs 96
RDY signal	IIGs 68	SHADOW (soft switch)	IIGs 70
READ (BASIC.SYSTEM)	PDOS 24	shadowing, super-hires	IIGs 72
READ (firmware routine)	IIE 10	sheet feeder	ATLK 9
ReadAsciiTime	IIGs 24	Shell	SC 1
ReadMouse	IIE 1, IIGs 85, Mous 1, 3	shell application, GS/OS	FTN \$B5
RealFreeMem	IIGs 51	shell programs and resources	IIGs 83
rebooting	IIGs 49	ShieldCursor	IIGs 34
recharge routine	IIGs 16	Shift key	IIGs 58
rectangle list resource format	IIGs 76	shift key modification	IIE 9, 10
REDIRECT, shell call	IIGs 30	Sholes keyboard	IIE 4
refCon, list control	IIGs 38	SHORTGRAPHICS unit	Pasc 15
regions	IIGs 10, 72	ShowPen	IIGs 72
registers, SCC	IIGs 18	ShrinkIt	Misc 14
Remote Print Manager	ATLK 3, 4, IIGs 77, 93	SIGNAL	GSOS 6
RemoveFromRunQ	IIGs 24	SizeWindow	IIGs 4
RESET routine	IIGs 49	slot 3	PDOS 15
reset key	IIE 5	slot arbitration	GSOS 4, 6, IIGs 69
ResetHook	IIGs 25, SmPt 6	slot mapping (ProDOS devices)	PDOS 3
ResetMember2	IIGs 24	slot ROM space	ATLK 4
resolution, changing	IIGs 4	SLOT3ROM (soft switch)	GSOS 2
Resource File	IIGs \$E/0001, PDOS 25	SLOTMSBL soft switch	IIGs 69
Resource Manager	IIGs 53, 76, 83, 86	Small Computer Systems Interface	see SCSI
resource application	IIGs 3, 83	SmartPort errors	IIGs 25
resource application, current	IIGs 83	SmartPort	IIE 10, IIGs 25, 26, PDOS 20, 21, SmPt all
resource converters	IIGs 86	SmartPort version	IIGs 100
resource files	GSOS 13	soft switches	IIE 10, Misc 15
resource forks	GSOS 5, IIGs 71		

disk port
 Sound Tools
 sound
 multi-channel
 resource format
 sampled
 sound clicks
 sound connector
 sound RAM
 sound scrap
 sound, compressed
 SoundEx
 SoundSmith
 source code file
 SPCCommand
 special memory
 spExtra
 SPOutStatus
 spreadsheet, AppleWorks
 SPWrite
 stack pointer
 stack space, CDAs
 stack, page one
 Standard Apple Numerics Environment
 Standard File
 StartFrameDrawing
 StartInfoDrawing
 startup order
 StartUpTools
 STATUS bug in UniDisk 3.5
 STATUS
 SmartPort
 status
 status string
 status, Pascal protocol
 statusBits
 stereo sound
 Sticky Keys
 storage types
 STORE (BASIC.SYSTEM)
 STROBE (signal)
 subtype, SmartPort
 Super Serial Card
 super hi-res graphics
 super hi-res screen
 super hi-res screen location
 super-hires screen image
 SuperInfo modules
 SW0 (soft switch)
 SW1 (soft switch)
 SW2 (soft switch)
 swap mode (sound)
 SysFailMgr
 system Monitor
 system software
 SYSTEM.APFLE
 SystemEvent
 SystemTask
 tables of addresses
 Tangent, Mr.
 TaskMaster
 task mask

 TaskMasterDA
 Teach document
 Teach
 TEClick
 TEFormat
 TEFormat structure
 TEGetRuler
 TEGetText
 TEIdle
 TEInsert
 TEKey
 telecommunications
 TEPaintText
 TERuler
 TETextRuler
 TETextText
 TestControl
 Text Edit controls
 Text Edit
 Text Tools
 text buffer, QuickDraw II
 text color
 text cursor bug
 text printing
 TextEdit style scrap
 TimeData
 timing
 65816 cycle
 65816 instruction
 titles, window and menu
 TLMountVolume
 TLShutDown
 TLStartUp
 TLTextMountVolume
 ToBaysStrip, vector
 Tool Locator
 tool dispatcher, patching
 tool patching
 tool set interdependencies
 tool set pointer table (TPT)
 tool sets
 user
 tool sets versions, system
 Toolbox Reference updates
 toolbox
 and desk accessories
 memory usage
 TOPRINTMSG
 ToStrip, vector
 TOWRITEBRAM
 TrackControl
 TrackGoAway
 transparent network printing
 TRESEARCH routine
 trust and verification
 TURTLEGRAPHICS unit
 type-ahead buffer
 typeless files
 UCSD Pascal
 UNBIND_INT_VECTOR
 UnBindInt
 UniDisk 3.5 drive
 UniDisk 3.5
 UniDiskStat
 UnionRgn
 unit_number

IIGs 30
 IIGs 26, 37, SC 10, 23

 IIGs 19
 IIGs 76
 FTN \$D8
 IIGs 11
 IIGs 9
 see DOC RAM
 IIGs 99
 FTN \$D8/0001
 SC 23
 FTN \$D5/0007
 FTN \$B0
 ATLK 5
 IIGs 52, 53
 IIGs 72
 ATLK 1
 FTN \$1B
 ATLK 5
 ATLK 2, IIGs 25, 70, 88
 IIGs 71
 IIGs 71, 88
 see SAME
 IIGs 96, SC 18
 SC 5
 IIGs 3
 IIGs 12
 IIGs 100, 12, 24, 53, 83
 UDsk 3
 UDsk 3
 SmPt 7
 PDOS 20, SmPt 2, 9, UDsk 2, 4
 GSOS 13
 ATLK 9
 ATLK 6
 ATLK 9
 IIGs 19
 IIGs 91
 PDOS 25
 PDOS 24
 IIE 10
 SmPt 7
 IIE 7, Misc 3
 Misc 17
 FTN \$B3, \$B5
 IIGs 91
 FTN \$C1/0000
 FTN \$C/4002
 IIE 9
 IIE 9
 IIGs 11
 IIGs 33
 IIGs 25
 GSOS 1
 Fasn 17
 IIGs 24
 IIGs 53, 91
 IIGs 90
 IIGs 71
 IIGs 42
 IIGs 84, 47, 53, 75, 80, 81, 84, 91,
 SC 2, 9
 IIGs 84
 FTN \$50/5445
 SC 24
 IIGs 92
 FTN \$50/5445
 IIGs 99
 IIGs 92
 IIGs 92
 IIGs 84
 IIGs 92
 IIGs 84
 Misc 14
 IIGs 92
 IIGs 92
 IIGs 92
 IIGs 92
 IIGs 81
 FTN \$50/5445, IIGs 92, SC 24
 GSOS 4, IIGs 26, 69, 71, 93
 IIGs 4
 IIGs 63
 IIGs 65
 IIGs 93
 IIGs 99
 Mous 2

 IIGs 70
 IIGs 2
 IIGs 5
 IIGs 24
 IIGs 53, 73, 101
 IIGs 53
 IIGs 24
 IIGs 73
 IIGs 26, 89
 IIGs 87
 IIGs 101
 IIGs 12
 IIGs 101
 FTN \$8A
 IIGs 73, SC 20
 IIGs 100
 IIGs 24

 IIGs 53
 IIGs 51
 IIGs 25
 IIGs 73
 IIGs 25
 IIGs 84
 IIGs 42
 ATLK 4
 Pasc 14
 FTN \$B8
 Pasc 15
 IIC 6
 FTN \$00
 Pasc all
 GSOS 9
 GSOS 9
 PDOS 23
 UDsk all
 IIGs 25
 IIGs 24
 PDOS 20, 21

UnitStatus (Pascal)
 UnitStatus
 Universal Access
 unknown files
 UnloadOneTool
 UnloadSeg
 UnloadSegNum
 UnPackBytes

 unpacked picture
 UnshieldCursor
 update routines, dialogs
 updateRgn
 updating windows
 User ID
 user names
 user tool sets
 UserCtlItem
 UserInfo
 UserItem
 userItems in dialogs
 VBL (softswitch)
 VBL interrupts
 rate
 VBL signal
 VCB
 vendor ID, SmartPort
 version
 version number, SmartPort
 version, resource type
 versions
 vertical blanking
 Video Keyboard
 Video Overlay Card
 video counters
 Visit Monitor desk accessory
 visRgn
 Volume Control Block
 volume control
 volumes (Pascal)
 WAIT (firmware routine)
 WAIT routine
 wait mode
 wFrame
 WindNewRes
 Window Manager
 window background patterns
 window defProc
 window information bar
 window record
 window titles
 windows, color tables
 windows, custom
 windows, updating
 wInfoDefProc
 wInfoHeight
 wInSpecial (TaskMaster result)
 word processor
 AppleWorks
 AppleWorks GS
 Medley
 Text Edit
 WordPerfect
 WordPerfect file
 work area pointer
 WRCARDRAM (soft switch)
 WRITE (firmware routine)
 WRITE (SmartPort)
 write-protect bug (UniDisk 3.5)
 XCMD/XPCN (HyperCard IIGS)
 XorRgn
 years, ProDOS
 Z8530 serial chip
 zero page
 SmartPort use of
 zero-crossing byte

Pasc 17
 IIC 5
 IIGs 91
 FTN \$00
 IIGs 53
 IIGs 22
 IIGs 22
 FTN \$08/4000, \$08/4001,
 \$C0/0001, \$C0/0002, IIGs 94
 IIGs 27
 IIGs 34
 IIGs 91
 IIGs 75
 SC 7
 IIGs 66, 71, PDOS 27
 IIGs 77
 IIGs 53, 73, SC 20
 IIGs 38
 ATLK 8, SC 18
 IIGs 38
 IIGs 91
 IIE 10
 IIGs 39, 40, Mous 2
 Mous 2
 IIC 9, IIGs 40
 PDOS 23, 8
 SmPt 2
 IIGs 76
 SmPt 2
 IIGs 76
 IIGs 100
 see VBL
 IIGs 91
 Misc 17
 IIGs 70
 IIGs 25
 IIGs 75, 80, 91
 see VCB
 IIC 7
 Pasc 16
 IIE 10
 Misc 12
 GSOS 13
 IIGs 3, 42
 IIGs 4
 IIGs 26, 47, 61, 62, 75
 IIGs 62
 IIGs 42
 IIGs 24, 3
 IIGs 42
 IIGs 5, 61
 IIGs 98
 IIGs 42, SC 5
 IIGs 83, SC 7
 IIGs 3
 IIGs 3
 IIGs 24

 FTN \$1A
 FTN \$50/8010
 FTN \$54/DD3E
 FTN \$50/5445
 FTN \$A0/0000
 FTN \$A0/0000
 IIGs 73
 IIGs 70
 IIE 10
 UDsk 4
 UDsk 3
 IIGs 86, HCGS 1
 IIGs 24
 PDOS 28
 IIGs 18
 PDOS 22
 SmPt 6
 IIGs 1

#0: About Apple II Technical Notes

March-July 1991

Technical Note #0 (this document) accompanies each release of Apple II Technical Notes. This combination release includes a new Notes for the Apple IIe #10, Apple IIgs #101, HyperCard IIgs #1, #2 and #3, and revised Notes for the IIgs #24, #43, #44, #53, #63, #71, #73, #76, #81, #83, #86, #90, #91 and 99, Miscellaneous #2 and #7 as well as an index to all released Apple II Technical Notes, File Type Notes, and Sample Code. If there are any subjects which you would like to see treated in a Technical Note (or if you have any questions about existing Technical Notes), please contact us at one of the following addresses:

Apple II Technical Notes
Developer Technical Support
Apple Computer, Inc.
20525 Mariani Avenue, M/S 75-3T
Cupertino, CA 95014
AppleLink: AIIDTS
MCI Mail: AIIDTS (264-0103)
Internet: AIIDTS@AppleLink.Apple.com

We want Technical Notes to be distributed as widely as possible, so they are sent to all Partners and Associates at no charge; they are also posted on AppleLink in the Developer Services bulletin board and other electronic sources, including the Apple FTP site (IP 130.43.2.3). You can also order them through APDA. As an APDA customer, you have access to the tools and documentation necessary to develop Apple-compatible products. For more information about APDA, contact:

APDA
Apple Computer, Inc.
20525 Mariani Avenue, M/S 33-G
Cupertino, CA 95014
(800) 282-APDA or (800) 282-2732
Fax: (408) 562-3971
Telex: 171-576
AppleLink: APDA
Internet: APDA@AppleLink.Apple.com

We place no restrictions on copying Technical Notes, with the exception that you cannot resell them, so read, enjoy, and share. We hope Apple II Technical Notes will provide you with lots of valuable information while you are developing Apple II hardware and software. The following pages list all Apple II Technical Notes that have been released.

Developer Technical Support thanks Dave Ely and Eric Mueller for their assistance in creating these text file versions of Technical Notes.

Released Apple II Technical Notes

July 1991

New ***
Revised *R*

Apple IIc

1	Mouse Differences On IIe and IIc	11/88
2	40-Column and Double High-Resolution Graphics	11/88
3	Foreign Language Keyboard Layouts	11/88

4	Dvorak Keyboard Layout	11/88
5	Memory Expansion on the Apple IIc	11/88
6	Buffering Blues	11/88
7	Existing Versions	11/88
8	Single-Sided 3.5" Media and the Apple IIc Plus	5/89
9	Detecting VBL	11/90

Apple IIe

1	Overview of the Apple IIe	11/88
2	Hardware Protocol for Doing DMA	11/88
3	Double High-Resolution Graphics	11/88
4	RDY line	11/88
5	/INH line	11/88
6	The Apple II Paddle Circuits	11/88
7	Interfaces--Serial, Parallel, and IEEE-488	11/88
8	Known Anomalies of Enhanced IIe ROMs	11/88
9	Switch Input Changes	11/88
***	10 The Apple IIe Card for the Macintosh LC	07/91

Apple IIgs

1	How to Install Custom BRK and /NMI Handlers	11/88
2	Transforming I/O Subroutines for Use in "Native" Mode	11/88
3	Window Information Bar Use	1/91
4	Changing Graphics Modes in Mid-Application	1/91
5	Window and Menu Titles	11/90
6	QuickDraw II Pattern Data Structure	7/89
7	Halt Mechanism in IIgs SANE	11/88
8	Elms Functions in IIgs SANE	11/88
9	IIgs Sound Expansion Connector: Analog Input/Output Impedances	11/88
10	InvalRgn Twist	11/88
11	Ensoniq DOC Swap-Mode Anomaly	11/88
12	Tool Set Interdependencies	1/90
13	ROM 1.0 Modem Firmware Bug	11/88
14	Standard File Calls and GrafPort Records	11/88
15	InstallFont and Big Fonts	7/89
16	Notes on Background Printing	11/88
17	Application Memory Management and the MMStartUp User ID	11/88
18	Do-It-Yourself SCC Access	7/90
19	Multichannel Output with the Apple IIgs Note Synthesizer	11/88
20	Catalog of APW Language Numbers	3/90
21	DMA Compatibility for Expansion RAM	11/88
22	Proper Use of Dynamic Segments	9/90
23	Toolbox Use of DOC RAM	11/88
R	24 Apple IIgs Toolbox Reference Updates	7/91
25	Apple IIgs Firmware Reference Updates	9/90
26	ROM Revision Summary	9/89
27	Graphics Image File Formats	11/88
28	Interface Card Design Guidelines	11/88
29	Monochrome High-Resolution Mode	11/88
30	Apple IIgs Hardware Reference Updates	9/90
31	Redirecting Output in APW C	11/88
32	/INH Line Anomaly	11/88
33	ERRORDEATH Macro	11/88
34	Low-Level QuickDraw II Routines	1/91
35	Printer Driver Specifications	9/90
36	Port Driver Specifications	9/89
37	Free-Form Synthesizer Tips	11/88
38	List Controls in Dialog Boxes	9/90
39	Mega II Video Counters	7/89

	40	VBL Signal		7/89
	41	Font Family Numbers		11/90
	42	Custom Windows		11/88
R	43	Undocumented Feature of CalcMenuSize	obsolete	7/91
R	44	GetPenState and SetPenState Record Error	obsolete	7/91
	45	Parameters for GetFrameColor		9/89
	46	DrawPicture Data Format		11/88
	47	What SetDataSize Does		11/88
	48	All About AlertWindow	obsolete	11/90
	49	Rebooting (Really)		1/89
	50	Extended Serial Interface Error Handling		1/89
	51	How to Avoid Running Out of Memory		9/90
	52	Loading and Special Memory		7/89
R	53	Desk Accessories and Tools		3/91
	54	MIDI Drivers		11/90
	55	Avoiding ClrHeartBeat		7/89
	56	Managing Dynamic Segments	obsolete	11/90
	57	Preventing Memory Compacting and Purging		7/89
	58	Keyboard Modifiers Register Anomaly		7/89
	59	Do Not Create Zero-Length Text Scraps	obsolete	1/91
	60	Menu Manager Memorabilia		11/90
	61	Window Title Handles		7/89
	62	No Non-Solid Window Background Patterns		7/89
R	63	Master Color Values		5/91
	64	Apple IIgs Installer and Installer Scripts		9/89
	65	Control-^ is Harder Than It Looks		9/89
	66	ExpressLoad Philosophy		9/90
	67	LaserWriter Font Mapping		11/89
	68	Tips for I/O Expansion Slot Card Design		9/89
	69	The Ins and Outs of Slot Arbitration		5/90
	70	Fast Graphics Hints		9/89
R	71	DA Tips and Techniques		7/90
	72	QuickDraw II Quirks		11/90
R	73	Using User Tool Sets		7/91
	74	A Faster List Manager Draw Routine		11/89
	75	BeginUpdate Anomaly		1/90
R	76	Miscellaneous Resource Formats		7/90
	77	Print Manager & AppleTalk Configuration Files		1/90
	78	Bank Alignment and Memory Management		3/90
	79	Integer Math Data Types		5/90
	80	QuickDraw II Clipping		3/90
R	81	Extended Control Ecstasy		7/91
	82	Controlling the Control Manager	obsolete	11/90
R	83	Resource Manager Stuff		3/91
	84	TaskMaster Madness		7/90
	85	Moving the Mouse		7/90
R	86	Risking Resourceful Code		3/91
	87	Patching the Tool Dispatcher		9/90
	88	The Page One Stack in a 16-Bit World		9/90
	89	MessageByName--Catchy Messages		9/90
R	90	65816 Tips and Pitfalls		3/91
R	91	The Wonderful World of Universal Access		7/91
	92	Twisted Tales of TextEdit		11/90
	93	Compatible Printing		9/90
	94	Packing It In (and Out)		9/90
	95	ROM Diagnostic Errors		9/90
	96	Standard File Customization		11/90
	97	Picture Comments and Printing		11/90
	98	Aren't Windows A Pane		1/91
R	99	Supplemental Scrap Types		3/91
	100	VersionVille		1/91
***	101	Patching the Toolbox		5/91

	1	80-Column Screen Dump	11/88
R	2	Apple II Family Identification Routines 2.2	5/91
	3	Super Serial Card Firmware Bug	11/88
	4	AppleWorks Keys	obsolete 5/89
	5	AppleWorks File Formats	obsolete 5/89
	6	IWM Port Description	11/88
R	7	Apple II Family Identification	5/91
	8	Pascal 1.1 Firmware Protocol ID Bytes	11/88
	9	AppleSoft Real Variable Storage	11/88
	10	80-Column GetChar Routine	9/89
	11	Examining the \$C800 Space from AppleSoft	5/89
	12	Apple II Firmware WAIT Routine	11/88
	13		not used
	14	Guidelines for Telecommunication Programs	7/89
	15	Compatibility Across Apple II Models	1/90
	16	Apple II Parallel Interface Card Firmware	7/90
	17	Buried Treasures of the Video Overlay Card	9/90

AppleTalk

	1	Identifying AppleTalk	3/90
	2	ProDOS 8 Compatibility on the IIe and IIgs	11/88
	3	Avoiding Remote Printer Time-Outs	9/89
	4	Printing Through the Firmware	9/90
	5	SPCommand Calls and Error \$0702	7/89
	6	Apple IIe Workstation Card Anomalies	7/89
	7	MLIACTV Flag and the IIe Workstation Card	11/89
	8	Using the @ Prefix	9/90
	9	The PAP Status Buffer	11/90

HyperCard IIGS

***	1	Corrections to the Script Language Guide	3/91
***	2	Known HyperCard Bugs	3/91
***	3	Pitching Sampled Sounds	3/91

GS/OS

	1	Contents of System.Disk and System.Tools	1/91
	2	GS/OS and the 80-Column Firmware	11/88
	3	Pointers on Caching	11/88
R	4	A GS/OS State of Mind	3/91
	5	Resource Fork Formats	7/89
	6	Drivers and GS/OS Direct Page	11/90
	7	Behavior of SET_DISKSW	7/89
	8	Filenames With More Than CAPS and Numerals	7/89
	9	Interrupt Handling Anomalies	5/90
	10	How Applications Find Their Files	9/90
	11	About EraseDisk and Format	11/90
	12	All About Notify Procs	9/90
R	13	GS/OS Reference Update	3/91

ImageWriter

	1	Custom Font Selection	11/88
--	---	-----------------------	-------

Memory Expansion Card

	1	Questions and Answers	11/88
--	---	-----------------------	-------

Mouse

	1	Interrupt Environment with the Mouse	11/88
--	---	--------------------------------------	-------

2	Varying VBL Interrupt Rate	11/88
3	Mode Byte of the SetMouse Routine	11/88
4	Mouse Firmware Bug Affecting ServeMouse	11/88
5	Check on Mouse Firmware Card	11/90
6	MouseText Characters	1/89
7	Mouse Clamping	11/88

Pascal

4	Pascal Declarations and the Directory Structure of a Blocked Volume	11/88
10	Configuration and Use of the Apple II Pascal Run-Time Systems	11/88
12	Disk Formatter Routine	11/88
14	Apple Pascal 1.3 TREESEARCH and IDSEARCH	11/88
15	Apple II Pascal SHORTGRAPHICS Module	11/88
16	Driver to Have Two Volumes on One 3.5" Disk	11/88
17	SYSTEM.APPLE Patch V2.0	3/90

ProDOS 8

1	The GETLN Buffer and a ProDOS Clock Card	11/88
2	Porting DOS 3.3 Programs to ProDOS and BASIC.SYSTEM	11/88
3	Device Search, Identification, and Driver Conventions	11/88
4	I/O Redirection in DOS and ProDOS	11/88
5	ProDOS Block Device Formatting	11/88
6	Attaching External Commands to BASIC.SYSTEM	11/88
7	Starting and Quitting Interpreter Conventions	11/88
8	Dealing with /RAM	11/88
9	Buffer Management Using BASIC.SYSTEM	11/88
10	Installing Clock Driver Routines	11/88
11	The ProDOS 8 MACHID Byte	11/88
12	Interrupt Handling	11/88
13	Double High-Resolution Graphics Files	11/88
14	Selector and Dispatcher Conventions	11/88
15	How ProDOS 8 Treats Slot 3	11/88
16	How to Format a ProDOS Disk Device	11/88
17	Recursive ProDOS Catalog Routine	11/89
18	/RAM Memory Map	11/88
19	File Auxiliary Type Assignment	11/88
20	Mirrored Devices and SmartPort	11/88
21	Identifying ProDOS Devices	3/90
22	Don't Put Parameter Blocks on Zero Page	7/89
23	ProDOS 8 Changes and Minutia	9/90
24	BASIC.SYSTEM Revisions	9/90
25	Non-Standard Storage Types	7/89
26	Polite Use of Auxiliary Memory	1/90
27	Hybrid Applications	1/90
28	ProDOS Dates--2000 and Beyond	9/90
29	Clearing the Backup Needed Bit	9/90

SmartPort

1	SmartPort Introduction	11/88
2	SmartPort Calls Updated	9/89
3	SmartPort Bus Architecture	11/88
4	SmartPort Device Types	11/88
5	SCSI SmartPort Call Changes	11/90
6	Apple IIgs SmartPort Errata	11/90
7	SmartPort Subtype Codes	11/88
8	SmartPort Packets	5/89
9	Apple II SCSI Errata	7/90

UniDisk 3.5

1	UniDisk 3.5 Internals	11/88
2	UniDisk 3.5 ID Bytes	11/88
3	STATUS Call Bug	11/88
4	Accessing Macintosh Disks	11/88
5	Architectural Differences Between 3.5" Drives	11/88

Apple IIGS

#1: How to Install Custom BRK and /NMI Handlers

Revised by: Jim Mensch & Jim Merritt

November 1988

Written by: Jim Merritt

October 1986

This Technical Note discusses a method to install a custom debugger or debugging stub within the Apple IIGS system.

Introduction

This Technical Note discusses a particular method that you may use to install a custom debugger or debugging stub within the Apple IIGS system. The strategy and techniques described here should be of special interest to those who wish to operate the Apple IIGS as a slave to a debugger that resides on another machine.

Typically, an interrupt handler should pass control to a debugger or debugging stub whenever the processor executes a BRK instruction, or when an interface card triggers a non-maskable interrupt (/NMI). To simplify the design of the debugger, the Apple IIGS Monitor should be responsible for the following:

- o saving all machine state information in locations that the debugger can access
- o setting the machine to a known state
- o passing control to an arbitrary debugger
- o restoring the remembered machine state upon regaining control from the debugger
- o resurrecting the interrupted process

The Monitor is designed to provide all of the services above for the BRK instruction, but only the third for /NMI interrupts. In addition, Apple II family systems are generally intolerant of /NMI interrupts. In this Technical Note we concentrate on the means by which you can install your own custom BRK handler, although we also briefly examine /NMI considerations.

Dealing With BRK

A BRK interrupt handler may reside at any address in memory. The Monitor passes control to your code by executing a JSL instruction; consequently, your routine must terminate with an RTL instruction. To install your BRK handler, simply load it into memory, call the Miscellaneous Tool Set GetVector routine to fetch the address of the current BRK handler, put that address in a safe place, then supply the address of your handler to the Miscellaneous Tool Set SetVector routine. To deactivate your handler, restore the previous handler address using SetVector as follows:

```
;  
; NOTE: All Listings are in APW assembler format.  
;
```

```
INSTMYBRK    anop                ;Example code to install user's BREAK handler.  
              PushLong #0        ;Space for function call result.  
              PushWord #$1C      ;We want BREAK vector address.
```

```

        _GetVector                ;Make the call using standard macro.

; The stack now holds address of the current break handler.
        PLA                        ;Get and save low word of address...
        STA      SBRKADR
        PLA                        ; ...and now high word.
        STA      SBRKADR+2
        PushWord #$1C              ;We want to change BREAK vector address.
        PushLong #MYHANDLR        ;Address of user's BRK handler.
        _SetVector                ;Make the call using standard macro.

; Custom handler is in place, now go off and do whatever we like...

DEACMYBRK    anop                  ;Example code to deactivate the BRK handler.
              PushWord #$1C        ;We want to change BREAK vector address.
              PushLong SBRKADR     ;The previous BRK handler address.
              _SetVector           ;Make the call using standard macro.

```

Upon entry to your code, the machine will be in eight-bit native mode. Specifically, the m and x bits will be set (forcing eight-bit accumulator, memory access, and index registers), the processor will be running at the normal (1 MHz) speed, all memory shadowing will be enabled, and both the direct page and data bank registers will be reset to zero. The same conditions must hold when your BRK handler returns control to the Monitor. While your code is active, however, it is free to affect the machine state in arbitrary ways, including (but not limited to) widening the registers, increasing the clock rate, and disabling shadowing. Before returning control to the Monitor, your break handler must also clear the processor's carry flag, as an indication that the BRK was indeed serviced by an external handler. (Note: The default BREAKVECTOR points to a "no-op" handler that simply sets the carry flag to indicate that there is no external handler available, and it then executes an RTL.)

When a BRK occurs, the processor saves the machine's state in the BRK.VAR area, and you may obtain this address with the Miscellaneous Tool Set GetAddr routine as follows:

```

        PushLong #0                ; space for result
        PushWord #9                ; we want BRK.VAR address
        _GetAddr                   ; make the call using standard macro

```

; The stack now holds the address of the BRK.VAR area, expressed as a long word (four bytes).

Coping With /NMI

Handling /NMI interrupts is, by far, a trickier proposition than fielding BRK instructions. For example, the user-definable /NMI jump-vector, /NMI (\$0003FB), only has room in its three-byte JMP-absolute instruction for a two-byte address. Because of this size limitation, at least the "front end" of any /NMI handler must reside in bank \$00. In addition, the Monitor does not "condition" the system in any way before transferring control through the /NMI hook, so the system could be in native mode, emulation mode, or any hybrid mode (with any screen condition) upon entry to your handler. (Note: Although the 65816 processor provides for separate /NMI vector addresses in native and emulation modes, the Apple IIGS implementation of these two vectors pass control to the same user hook at \$0003FB.) The processor only saves minimal machine state information when an /NMI occurs; if the handler needs to preserve more than the program counter and status register (which are saved automatically), then it must do so explicitly. Because the 65816 assumes any program running in emulation mode has its program bank register in bank zero,

it will not save the program bank register for any program running in emulation mode outside of bank zero. Code which runs in this manner will always crash if it makes any attempt to return from the interrupt. Finally, /NMI interrupts can create havoc with disk access and other aspects of the system; consequently, the only way you can safely use /NMI interrupts is as a one-way "escape hatch" to emergency debugging code.

Here are some ground rules for /NMI interrupt handlers.

- o On entry, store any interesting registers or machine state in RAM space owned by the handler.
- o Determine whether the processor is in emulation mode or native mode.
- o Take appropriate action, depending upon the processor mode.
- o Under no circumstances try to return from the interrupt! Restart the system instead.

To install an /NMI handler, load it into some free RAM in bank \$00, put the two-byte address currently at location /NMI+1 in a safe place, then replace it with the address of your handler. To deactivate your handler (assuming nothing has yet invoked it), simply restore the previous handler address to /NMI+1.

Apple IIGS

#2: Transforming I/O Subroutines for Use in "Native" Mode

Revised by: Pete McDonald
Written by: Pete McDonald

November 1988
October 1986

This Technical Note outlines a number of techniques useful when transforming Apple II I/O subroutines for use in the "native" Apple IIGS environment.

The Apple IIGS execution environment represents quite a departure from the environment to which the average Apple II developer is accustomed. This fact results in a number of unique problems when one attempts to convert existing Apple II applications for use in the "native" Apple IIGS environment. (Note: If you intend to let your application remain an eight-bit "classic" Apple II application, then you can ignore the information this Technical Note presents.)

I/O subroutines which depend upon critically timed code present some of the biggest conversion problems due to two major issues. In the native IIGS environment, you cannot guarantee that there will be memory available in a given bank, and I/O locations are not available in every bank.

There are a number of possible solutions to this problem. Which ones you should use depend upon what the program in question is doing. This Note attempts to describe some of the problem situations and possible solutions.

Examine the 6502 code segment below. It serves no useful purpose, other than to illustrate a simple manifestation of the problem. Assume IoLoc is a location in the \$C000 - \$CFFF range of memory.

```
Loop   LDA   IoLoc
        DEY
        BPL   Loop
```

Because the \$C000 - \$CFFF range of memory in bank 2 or higher contains RAM instead of I/O circuitry unless hardware shadowing is enabled, if you place the fragment above in one of these banks, it will have no effect on the I/O device you intend it to control.

There are two possible solutions in this case. Either change the instruction LDA IoLoc so it uses long addressing, thereby forcing the CPU to reference the proper bank. (Note: The problem with this is the long version of LDA requires an extra CPU cycle to execute. If the code segment is timing critical, then this method is likely to be unacceptable.) Alternately, in the timing-critical case, we could set the data bank register before entering the loop which would mean the LDA IoLoc would take the same number of cycles as it did previously, thus leaving the timing loop unchanged.

These solutions seem pretty easy; therefore, you know there is a catch. The catch, unfortunately, is that most code is not isolated as in the example. Specifically, code commonly tries to load from or store to some location in memory other than the I/O location at the same time it is trying to access the I/O location.

Take, for example, the following fragment:

```

Loop   LDA   Data,y
        STA   IoLoc
        DEY
        BPL   Loop

```

In this example, we assume that the label Data refers to some kind of table which normally resides in the same bank as the program. Now if you set the data bank register to access I/O locations, then the reference to Data will also reference the same bank as the I/O; this solution is likely not acceptable. One thing you can do is move the data table to the direct page (zero page for 6502 programmers), but now the LDA Data,y instruction will take one less cycle to execute. There is a solution, although it is a little complicated. If we set the direct page register to a non page-aligned location, then we effectively apply a one-cycle penalty to all direct page references and solve our problem.

Of course, nothing is ever as simple as it seems. What happens to references to other direct page locations that expect to operate without the one-cycle penalty? To properly address this question, I would need much more space than I have here, so in lieu of further examples, I offer some general information. (As an aside, I used these techniques to transform the old "Apple II Disk II formatter module" for use in any bank of memory in the native IIGS environment. I accomplished this using, almost exclusively, editor find and replace commands, and I finished in hours instead of the days which would have been required to completely rewrite the program.)

In addition to the techniques already covered, there are a few other things which may be necessary to complete a transformation (they were necessary in the case of the formatter module).

As I already mentioned, one problem is what to do in the case where a program references I/O, local program-bank data, and the zero-page. In this case, significant rewrites could be required, but not necessarily.

In the case of the disk formatter, it turned out that some modules used both normal zero-page addressing and normal 16-bit absolute indexed addressing. Since the transformation process dictates that we change 16-bit absolute addressing to direct-page addressing with a non page-aligned direct page, there could have been a problem had both uses of the direct page been timing critical. Fortunately, by treating each module of the program separately, when I needed both types of addressing, only one was critical. The solution was to set the direct page to a non page-aligned value in some modules and to a page-aligned value in others. There are some minor logistical issues when a direct page's base address can be at either \$xxx0 or \$xxx1, the biggest of which is keeping track of which is in effect at a given point and knowing to reference the label as label, label+1, or label-1, depending upon the particular case.

With the formatter transformation, there was one other major issue: there are not direct-page versions of all the 16-bit absolute addressing modes (i.e., one cannot convert 16bitaddress,x to 8bitaddress,x). In the case of the formatter, I was able to solve this by reversing all the register use (i.e., all LDY instructions became LDX instructions, all STY instructions became STX instructions, etc.).

There are still a number of other ways in which one can approach these issues; one that comes to mind would be using some form of the new stack-relative addressing modes to yield yet another range of semi-independently accessible addresses.

The real point of this Technical Note is that with a little thought and

effort, one can successfully convert a large subset of likely configurations for use in the native IIGS environment without major rewrites. The bottom line is to be creative!

Further Reference

- o Programming the 65816 Including the 6502, 65C02, and 65802 (Eyes/Lichty)
- o Apple IIGS Firmware Reference

Apple IIGS
#3: Window Information Bar Use

Revised by: Dave Lyons
Written by: Dan Oliver

January 1991
October 1986

This Technical Note details the use of a window's information bar, including a code sample which places a menu in an information bar.

Changes since November 1988: Added a note about the current Resource Application when inside an InfoDefProc procedure, and information about information bars and NewWindow2.

Apple IIGS window information bars are not as straightforward as other window features, and one reason for this is the small amount of space originally allocated for their processing. If you feel your application can benefit from the use of information bars, you can implement them, and this Technical Note explains how to do it and includes some suggestions for their use. The code samples below demonstrate how to place a menu bar in an information bar, but your use of information bars is not limited to those described here.

Information Bar Initialization

You can create an information bar in a window when you create the window by setting the following fields in the parameter list you pass to NewWindow:

wFrame	Set bit 4.
wInfoHeight	Set to the height of the information bar (should not exceed window height).
wInfoDefProc	Set to the address of the information bar definition procedure (see below).

If you create a window as visible, the Window Manager will call your information bar definition procedure (InfoDefProc) before returning from NewWindow. If you have to create the contents of the information bar after the window, you will have a problem since the Window Manager will expect your InfoDefProc to draw things which do not yet exist. You can solve this problem by creating the window as invisible, creating the contents of the information bar, then showing the window. Another solution would be to detect, in the InfoDefProc, that the contents of the information bar do not yet exist.

NewWindow2, however, does not let you override the information bar drawing procedure in the template. If you pass a window template in a resource, creating the window as visible crashes (since the address of your information bar drawing procedure cannot possibly be in the window template resource). Instead, create the window as invisible and call SetInfoDraw to set the address of the information bar drawing procedure before calling ShowWindow.

Below is an example of initializing a window's information bar to contain a menu bar. The three key fields of the parameter list which you pass to NewWindow are as follows:

wFrame Set bit 4 = 1 and bit 5 = 0 for an invisible window; the other bits do not affect the information bar, so you can set them as you wish.

wInfoHeight Assuming you are using a system menu bar and initializing it before the window, set to the height FixMenuBar returned when you created the system menu bar. If you would rather use an absolute value, which we do not advise, you could use 14 which should be about right for the current system font.

wInfoDefProc Set to the address of the InfoDefProc, in this case draw_info.

After you create the window, but before you show it, you can create the menu bar to place in the information bar. The code to create the menu bar might look like the following:

```

window            Direct page location that contains pointer to window's port.
;
; --- Create a menu bar
-----
;
;            pha                            Space for result.
;            pha
;            pea     $FFFF                 Set "use current port" flag.
;            pea     $FFFF
;            _NewMenuBar                 Create a menu bar.
;            pla                           Get returned menu bar handle.
;            sta     <menuBar             Remember menu bar handle.
;            pla
;            sta     <menuBar+2
;
;
; --- Store menu bar's handle in the window's InfoRefCon
-----
;
;            pei     <menuBar+2            Pass menu bar handle.
;            pei     <menuBar
;            pei     <window+2            Window to set refCon.
;            pei     <window
;            _SetInfoRefCon              Store menu bar handle in window's
infoRefCon.
;
;
; --- Make the window's menu bar the current menu bar
-----
;
;            pei     <menuBar+2            Pass menu bar handle.
;            pei     <menuBar
;            _SetMenuBar                 Make new menu bar the current menu bar.
;
;
; --- Get the RECT of the window's information bar
-----
;
;            pea     tempRect|-16         Pass pointer of RECT.
;            pea     tempRect
;            pei     <window+2            Pass pointer of window.
;            pei     <window
;            _GetRectInfo                tempRect = interior RECT of window's
Info Bar.

```

; --- Dereference menu bar handle

```
-----  
;  
    ldy    #2  
    lda    [menuBar],y  
    tay  
    lda    [menuBar]  
    sta    <menuBar  
Bar.    sty    <menuBar+2  
;  
;  
; --- Set size of menu bar  
-----
```

```
;  
;  
    lda    <tempRect+y1  
    dec    a  
    ldy    #CtlRect+y1  
    sta    [menuBar],y  
;  
    lda    <tempRect+x1  
    dec    a  
    ldy    #CtlRect+x1  
    sta    [menuBar],y  
;  
    lda    <rect+y2  
    inc    a  
    ldy    #CtlRect+y2  
    sta    [menuBar],y  
;  
;  
; --- Set flag to tell Menu Manager to draw menu in current port  
-----
```

```
;  
    ldy    #CtlOwner+2  
    lda    [menuBar],y  
    ora    #$8000  
    sta    [menuBar],y  
;  
;  
; --- Create the menus and add them to the window's menu bar  
-----
```

```
;  
loop    lda    #4  
        pha  
        tay  
;  
        pha  
        pha  
        lda    menu_list+2,y  
        pha  
        lda    menu_list,y  
        pha  
        _NewMenu  
;  
        pea    0  
        _InsertMenu  
;  
        pla
```

Now menuBar is the pointer to the Menu

Overlap top side.

Overlap left side.

Overlap bottom side.

Set high bit in CtlOwner.

Save index into menu list.
Switch index to Y.

Space for return value.

Pass address of menu/item lines.

Menu handle already on stack.
Insert menu list at front of list.
Add my menus to the system menu bar.

```

        sec
        sbc     #4
        bpl     loop
;
;
; --- Initialize the size of the menu bar and menus
-----
;
        pha                Space for returned bar height.
        _FixMenuBar        Fix up positions in the menu bar.
        pla                Discard height of menu bar.
;
;
; --- Restore the system menu bar as the current menu
-----
;
        pea     0                Pass flag for system menu bar.
        pea     0
        _SetMenuBar        Make system menu bar current.

```

The window's menu bar is now initialized, and you can make the window visible with a call to ShowWindow; the InfoDefProc will draw the menu bar.

Information Bar Definition Procedure (InfoDefProc)

The InfoDefProc is slightly misleading; it is only responsible for drawing the interior, above the background, of the information bar. The InfoDefProc is not responsible for defining the information bar, drawing the frame and background, testing for hits, or tracking the user. The InfoDefProc is located inside your application, and the Window Manager calls it whenever it needs to draw the part of the window frame that contains the information bar.

Each window with an information bar can have its own InfoDefProc, or they can all share a common InfoDefProc. When the Window Manager calls your InfoDefProc, it sets the proper port, the Window Manager's port, and the proper state, an origin local to the window frame and clipped to any windows above it. The direct page and data bank are not defined and should be considered unknown.

The Window Manager passes your InfoDefProc the following information:

- o Pointer to the information bar's interior rectangle (less frame), local coordinates.
- o Value of the window's wInfoRefCon, set and used only by your application.
- o Pointer to the window's port (do not switch to this port for drawing).

A window that has an information bar containing a menu bar (handle stored in the window's InfoRefCon) might have a InfoDefProc as follows:

```

draw_info    START
;
theWindow    equ    6                Offset to the information bar owner window
infoRefCon   equ    theWindow+4     Offset to the window's information bar RefCon
infoRect     equ    infoRefCon+4    Offset to the information bar's enclosing RECT
;
        phd                Save original direct page.
        tsc                Switch to direct page in stack.
        tcd
;
;
; --- Draw the window's menu bar in the window's information bar
-----
;

```

```

pei    infoRefCon+2    Pass handle of window's menu bar handle.
pei    infoRefCon
_SetMenuBar           Make the window's menu bar the current menu
                      bar.
;
  _DrawMenuBar        Draw the window's menu bar, as requested.
;
  lda    #0           Zero is the flag for the system menu bar.
  pha
  pha
  _SetMenuBar        Make the system menu bar current again.
;
;
; --- Remove input parameters from the stack
-----
;    ldx    #12
    ply           Pull original direct page, save in Y.
;
    tsc           Move direct page point to stack.
    tcd
    lda    2,s      Move return address over input parameters.
    sta    2,x
    lda    0,s
    sta    0,x
;
    tsc           Adjust stack for stripped input parameters.
    phx           Number of bytes of input parameters.
    clc
    adc    1,s      Add number of input parameters to stack
                   pointer.
    tcs           And reset stack.
;
    tya           Restore original direct page.
    tcd
;
    rtl           Return to Window Manager.
    END

```

Information Bar Environment

An information bar is part of a window's frame, that is, not part of the window's content region. Because it is part of the frame, an information bar is in the Window Manager's port, so before an interaction (drawing or mouse selecting), the proper port (Window Manager's) must be in the proper state. The proper state means the origin must be at the window's upper-left corner and clipped to any windows above.

When the Window Manager calls the InfoDefProc it sets the proper port to the proper state; however, to interact with the information bar outside the InfoDefProc, you must set the proper port the the proper state. You can accomplish this with a call to StartInfoDrawing. When the interaction is completed, you must allow the Window Manager to return its port to a general state via a call to EndInfoDrawing. You are in a special state that requires some constraints (discussed later) between the calls to StartInfoDrawing and EndInfoDrawing.

Here is an example of interacting with our window's menu bar.

```

;
poll   pha           Space for return value.
       pea    %0000111101101110    Pass event mask to use.

```

```

    pea    TaskRec|-16          Pass pointer to Task record.
    pea    TaskRec
    _TaskMaster
    pla
    beq    poll                Does event need further processing?
;
;
; --- Handle button down in window's information bar
-----
;
    cmp    #InInfo            In Information bar?
    bne    poll
;
    pha
    pha
    lda    TaskRec+TaskData+2  Pass pointer of window.
    pha
    lda    TaskRec+TaskData
    pha
    _GetInfoRefCon            Get menu bar handle from window's
                                InfoRefCon.
    pla
    sta    menuBar
    pla
    sta    menuBar+2
;
;
; --- Switch to proper port in proper coordinate system
-----
;
    pea    tempRect|-16       Pass pointer to RECT to store info
                                bar RECT.
    pea    tempRect
    lda    TaskRec+TaskData+2  Pass pointer of window.
    pha
    lda    TaskRec+TaskData
    pha
    _StartInfoDrawing
;
;
; --- Handle menu selection from window's menu bar
-----
;
    pea    TaskRec|-16       Pass pointer to Task record for
                                MenuSelect.
    pea    TaskRec
    pei    menuBar+2         Pass handle of menu bar.
    pei    menuBar
    _MenuSelect              Let user make selection.
;
    lda    event+TaskData     Get the item's ID number.
    beq    exit              Was a selection made?
;
    _EndInfoDrawing          Switch back to original port.
;
;
; (Handle the menu selection.)
;
; The EndInfoDrawing followed by the StartInfoDrawing call is only
; needed when code between them calls the Window Manager.
;
    pea    tempRect|-16       Pass pointer to RECT to store info

```

```

bar RECT.
    pea    tempRect
    lda    TaskRec+TaskData+2      Pass pointer of window.
    pha
    lda    TaskRec+TaskData
    pha
    _StartInfoDrawing             Switch to the proper port in the
proper state.
;
    pea    0                       Pass unhilite flag.
    lda    TaskRec+TaskData+2     Pass menu's ID number.
    pha
    _HiliteMenu                   Unhilite menu's title.
;
;
; --- Clean up and return to polling
-----
;
exit    _EndInfoDrawing           Switch back to original port.
;
    pea    0                       Make system menu bar current.
    pea    0
    _SetMenuBar
;
    jmp    poll                   Return to polling user.
;

```

Information Bar Shutdown

When the Window Manager closes the window, it is up to you to resolve any shutdown necessities associated with the information bar. Using our window menu bar example, the close window might look like the following:

```

;
    pei    menuBar+2              Pass handle of menu bar
    pei    menuBar
    _SetMenuBar
;
    pha                            Space for returned menu handle.
    pha
    pea    2                      ID number of second menu.
    _GetMHandle                   Get the menu's handle.
    _DisposeMenu                 Free menu record and associated data.
;
    pha                            Space for returned menu handle.
    pha
    pea    1                      ID number of first menu.
    _GetMHandle                   Get the menu's handle.
    _DisposeMenu                 Free menu record and associated data.
;
    pea    0                      Make system menu bar current.
    pea    0
    _SetMenuBar
;
    pha                            Space for menu bar's handle.
    pha
    pei    <window+2              Pass pointer of window to close.
    pei    <window
    _GetInfoRefCon               Get the InfoRefCon from the window.
    _DisposeHandle               Free menu bar record.
;

```

```
pei    <window+2           Pass pointer of window to close.
pei    <window
_CloseWindow             Now the window can be closed.
```

;

The type of shutdown you use depends upon the contents of the informationbar.

Why didn't I put a DisposeMenuBar call in the Menu Manager? I didn't think of it until a week too late. Sorry.

Other Information Bar Uses

The following suggestions are only theories and have not been tested.

- o Display text information, as in Macintosh Finder windows.
- o Split window. Like the content region, the information bar could be large enough to hold data.
- o Hold controls. You could scroll data in the content region while keeping the controls which affect the display in place and within the user's reach. (Note: The Control Manager currently will not allow controls it creates in an information bar. In this case, NewControl would be using a port that is not in your window's port, namely the Window Manager's port.)

Further Reference

-
- o Apple IIGS Toolbox Reference, Volumes 1 & 2

Apple IIgs

#4: Changing Graphics Modes in Mid-Application

Revised by: Dave "Dave" Lyons, C.K. Haun & Dan Oliver

January 1991

Written by: Dan Oliver

October 1986

This Technical Note discusses how to switch between the two graphics modes, 320 and 640 horizontal resolution, while running an application which uses the Window, Control, and Menu Managers.

Changes since May 1990: Added information about reinstalling fonts after restarting QuickDraw II.

Why Change Resolution?

Why not? There are certain applications where the ability to run in both modes is essential; most graphics applications fall into this category. Other applications might switch modes to provide features which their competitors lack; a financial application might display figures in 640 mode and charts in 320 mode. Still other applications may want to give the user the choice. A word processor might seem useful only in 640 mode, but what if the user wants to print greeting cards with pictures? The user does not need the line length provided in 640 mode but does need the added color of 320 mode for the pictures.

Let me preach a little. I have worked on other machines with different graphic modes and learned some things that might be of use to application programmers. Many application programmers fight mode switching with either rhetoric or apathy, then when users expect their software to run in either mode, they become frustrated when it does not allow switching. To avoid the problem of frustrating the user, you can provide mode switching (which is not as hard as you might think).

How To Change Modes

First, assume you are in an application which is running with a system menu bar, a few visible windows with scroll bars, and one window with some standard controls. At some point, the user decides to change modes, possibly via a menu item thoughtfully provided by the application programmer. Your change mode handler might look like the following:

```
;  
; --- This step is necessary if QuickDraw Auxiliary is started -----  
;      _QDAuxShutDown      ;Shut down QDAux first  
; -----  
;      _QDShutdown         ;Shut down QuickDraw.  
;                          ;This will turn graphics off so you will see  
;                          ;the text screen for a second (a advertisement  
;                          ;might go here).  
;      lda <mode           ;Variable that holds current resolution.  
;      eor #$0080         ;Flip the mode bit, $0000 = 320, $0080 = 640.  
;      sta <mode          ;New value will be used to start the new mode  
;  
;      pei <QDzpage        ;Pass the direct pages allocated for  
;                          ;QuickDraw.  
;      pei <mode          ;New mode.  
;      pei <QDwidth       ;0 for screen width; other numbers for
```

```

                                ;printing
pei <MyID                       ;Pass my ID number.
_QDStartup                       ;Restart QuickDraw in the new mode.
;
    _GrafOff                     ;Turn screen off because changing mode
                                ;may not be pretty.
; --- This step is necessary if you need QuickDraw Auxiliary -----
    _QDAuxStartUp               ;Start QDAux again
; -----
;
; --- Fix up the cursor for the new mode -----
;
    pea 0                       ;Pass minimum cursor X position.
    lda #319                    ;Maximum X position for 320 mode.
    ldx <mode                   ;320 or 640 mode?
    beq store
    lda #639                    ;Maximum X position for 640 mode.
store pha                       ;Pass maximum cursor X position.
    pea 0                       ;Pass minimum Y cursor position.
    pea 199                     ;Pass maximum Y cursor position.
    _ClampMouse                 ;Clamp the cursor to the new screen size.
;
    _HomeMouse                  ;Move the cursor to 0,0 to make sure
                                ;it is on screen.
    _ShowCursor                 ;Make cursor visible.
;
; --- Tell tools about the change -----
;
    _WindNewRes                 ;Tell Window Manager about the change.
    _MenuNewRes                 ;Tell Menu Manager about the change.
    _CtlNewRes                  ;Tell Control Manager about the change.
;
; --- Fix the screen to look good -----
;
    Here you might want to change the color of the desktop, windows, menus ;
or controls to look good for the new mode.
;
    See example below.
; --- Redraw the screen in the new mode -----
;
    pea 0                       ;Pass flag to draw entire screen.
    pea 0
    _RefreshDesktop             ;Draw entire screen.
;
    _GrafOn                     ;Now show the new screen.
;

```

That is not too bad, but I left out the fun part. Before the RefreshDesktop there is a section named "Fix up the screen to look good." This section is where you might want to put some color into windows, controls, and menus if you are switching to 320 mode; changing colors is not required, but there are some things which are.

When switching from 640 mode to 320 mode, some windows (both visible and invisible) might be positioned off the screen in 320 mode. The first way to handle this problem is easy for you, the programmer, but not so great for the user: close all the windows before changing modes, then position them correctly when the user opens them in the new mode. The second way to handle

the problem is to walk the window list and move all the windows, maybe even change their sizes. You could double each window's horizontal starting position and width when switching from 320 mode to 640 mode and halve it when changing from 640 mode to 320 mode. The vertical position and height are okay. An example of the second method is given below.

Windows with vertical scroll bars in the window frame are the same width when you change modes, so switching from 320 mode to 640 mode results in a narrower bar while changing from 640 mode to 320 mode produces a wider bar. The bars change to the correct size as soon as the user resizes the window, since `SizeWindow` deletes the old scroll bars and allocates new ones according to the current mode. If, as suggested above, you resize all the windows after the mode change and before calling `RefreshDesktop`, you should be in good shape. If you choose not to follow this recommendation, you should call `SizeWindow` for every window with scroll bars and change the size of each window at least one pixel since `SizeWindow` does not do anything if the passed size is not different than the current size.

You should dispose of scroll bars in a window's content region and recreate them; this is not nice, but very few applications have scroll bars in a window's content region.

You should not resize any open new desk accessory (NDA) windows. NDAs may be dependent on screen mode, or their current position, or other such things which may change with resolution. To be kind to the NDAs, you should issue a `CloseAllNDAs` call. This call allows the NDAs to go through their normal close procedures. If a user wants an NDA open in the new screen resolution he must reopen it. This assures that the NDA always knows its own position and the current screen resolution.

`WindNewRes` resets the desktop shape and pattern and the Window Manager's icon font to their defaults for the new mode, so if you changed any of these, you must add to or subtract from the desktop again and reinitialize to your custom pattern or icon font again.

`CtlNewRes` resets the Control Manager's icon font to the default for the new mode, so if you changed the Control Manager's icon font, you must reinitialize to your icon font again.

Reinstalling Large Fonts

After restarting `QuickDraw II`, you should call `InstallFont` again on the fonts your application is using. This causes the Font Manager to call `InflateTextBuffer` so that `QuickDraw` can draw text correctly in large font sizes.

Repositioning and Resizing Windows in the New Mode

Here is an example of how to reposition and resize windows in the new mode.

```
; QuickDraw and the tools have already been reinitialized in the new mode.
; mode = $0000 if in 320 mode, $0080 if in 640 mode.
;
BoundsRect    equ 8           ;Offsets in port record from QuickDraw document
PortRect      equ 16
;
                _CloseAllNDAs    ; close all open NDA windows
                pha              ;Space for result.
                pha
                _FrontWindow     ;Start with the top most window, this assumes
```

```

bra enter          ;there are no invisible windows ahead of the
                  ;active window in the window list.

ldy #BoundsRect+2
lda [window],y    ;Get window's starting horizontal position.
eor #$FFFF        ;Convert to screen coordinate (negate it).
inc a
asl a             ;Double it if we're going to 640 mode.
ldx <mode         ;Going to 320 or 640 mode?
bne store1       ;Ready if we're going to 640.
lsr a            ;Otherwise, undo the doubling,
lsr a            ;and halve the starting horizontal position.

store1            pha          ;Pass window's new X starting position.
                  ldy #BoundsRect
                  lda [window],y    ;Get window's starting vertical position.
                  eor #$FFFF        ;Convert to screen coordinate.
                  inc a
                  pha          ;Pass window's current Y starting position.
                  pei <window+2     ;Pass window to move.
                  pei <window
                  _MoveWindow      ;Move the window to its new position.

;

ldy #PortRect+6   ;Get window's current width.
                  lda [window],y    ;(This assumes the window's origin is 0,0.)
                  asl a             ;Double the window's width if going to 640 mode
                  ldx <mode         ;Going to 320 or 640 mode?
                  bne store2       ;Ready if we're going to 640.
                  lsr a            ;Otherwise, undo the doubling,
                  lsr a            ;and halve the window's width.

store2            pha          ;Pass window's new width.
                  ldy #PortRect+4
                  lda [window],y    ;Get window's height.
                  pha          ;Pass window's current height.
                  pei <window+2     ;Pass window to resize.
                  pei <window
                  _SizeWindow      ;Resize the window.

;

pha              ;Space for result.
pha
pei <window+2    ;Pass pointer to window we just processed.
pei <window
_GetNextWindow  ;Get the pointer to the next window.

;
enter           pla          ;Remember the pointer to this window.
               sta <window
               pla
               sta <window+2

;

ora <window     ;Are there any more windows?
bne loop

;

```

WindNewRes

Generally, WindNewRes does the following:

- o closes its port
- o opens its port again, now in the new mode
- o reinitializes the desktop size
- o chooses the proper icon font for close and zoom boxes
- o reinitializes the desktop pattern
- o changes the SCB byte of each window's port to the new mode

- o recomputes the VisRgn for each window

MenuNewRes

Generally, MenuNewRes does the following:

- o closes its port
- o opens its port again, now in the new mode
- o reinitializes internal parameters, like vertical line width, for the new mode
- o reinitializes the color palette via InitPalette
- o subtracts the system menu bar from the desktop (this is why you must call WindNewRes first)
- o draws the system menu bar

CtlNewRes

Generally, CtlNewRes does the following:

- o chooses the proper icon font for radio button, check box, grow box and scroll bar arrows
- o reinitializes internal parameters, like vertical line width, for the new mode

Apple IIGS
#5: Window and Menu Titles

Revised by: Matt Deatherage
Written by: Dan Oliver

November 1990
October 1986

This Technical Note discusses spacing for both window and menu titles. Changes since November 1988: Revised to include new information on the default placement of the Apple menu.

Strings used for window titles should always have a space as the first and last characters. This spacing is especially important for windows that use a lined window title bar since, without the beginning and ending space, the line pattern in the title bar runs against the title. Since there will be window editor desk accessories which allow the user to change the title bar pattern without the application knowing, you should pad your window titles with spaces even if you are using black window title bars.

The Window Manager does not force spaces on either side of titles to optimize the window frame drawing speed; it is much faster to let the text punch a hole in the title bar pattern than to compute the rectangle, fill it, and draw the text.

To provide the user with a consistent visual interface, you should also pad your menu titles with spaces. If you use either one or two spaces (the Apple IIGS Finder has used two) before and after each menu title, your menu titles will be consistent and balanced (two spaces work well in 640 mode where one space usually suffices for 320 mode). Although it is true that a menu bar will look about the same if the first menu title has two spaces before it and no space following it and all the other menu titles have four spaces before them, when the user pulls down the menu, the Menu Manager's highlighting will clearly (and embarrassingly) show the spaces in the menu titles.

If you would like to place the Apple menu differently, you must use Menu Manager calls since you cannot place spaces around the at sign (@) which the Menu Manager uses to represent the Apple logo in a menu title. The easiest way to accomplish this is calling `SetMTitleStart` to set the starting position for the leftmost title (usually the Apple menu) within the current menu bar. The Apple IIGS Finder has used a value of 10 (\$0A) pixels.

Beginning with System Software 5.0, the Apple menu is placed at a default of 10 pixels from the left edge of the menu bar in 640 mode or five pixels in 320 mode. If you use `SetMTitleStart` to change the default, the value is still interpreted as an absolute placement from the left edge of the menu bar. For example, `SetMTitleStart(6)` moves the Apple menu one pixel to the right of the default in 320 mode and four pixels to the left of the default in 640 mode. Be sure not to use `SetMTitleStart` to set the Apple menu starting place to the left of the default, as doing so interferes with the AppleShare activity arrows.

Apple IIGS
#6: QuickDraw II Pattern Data Structure

Revised by: Dave Lyons
Written by: Guillermo Ortiz

July 1989
December 1986

Some QuickDraw II calls require a pen pattern as input or return one as output; regardless of the drawing mode (320 mode or 640 mode), a pen pattern takes 32 bytes.

Changed since November 1988: Starting with System Software 5.0, all 32 bytes are significant if bit 15 of the current port's arcRot field is set. Changed wording to cover QuickDraw II patterns in general, instead of pen patterns only.

Early QuickDraw II documentation described the pattern data structure as follows:

TYPE

```
nibble = 0..15;
twobit = 0..3;
Pattern = RECORD CASE MODE OF
    mode320:(PACKED ARRAY [0..63] OF nibble);    { 32 bytes }
    mode640:(PACKED ARRAY [0..63] OF twobit);    { 16 bytes }
END;
```

This declaration could lead one to believe that 16 bytes are enough when making calls to QuickDraw II in 640 mode. This is not true. A pattern always takes 32 bytes; QuickDraw II calls that copy or construct patterns access all 32 bytes. That means it is never safe to pass the address of a 16-byte area as a pattern. Toolbox calls that return data into your buffer overwrite 16 bytes immediately following your buffer. Calls that copy data from your buffer access those extra 16 bytes, possibly including soft switches or reserved space in the memory map.

The difference between modes is that QuickDraw II normally ignores the second 16 bytes if the current port's locInfo indicates 640 mode. Starting with System Software 5.0, all 32 bytes of patterns are significant in 640 mode when bit 15 of the current port's arcRot field has been set with SetArcRot. In this case, patterns are 16 pixels wide and 8 pixels high.

Further Reference

- o Apple IIGS Toolbox Reference, Volume 2
- o System Software 5.0 documentation (APDA)

Apple IIGS

#7: Halt Mechanism in IIGS SANE

Revised by: Guillermo Ortiz & Matt Deatherage

November 1988

Written by: Guillermo Ortiz

December 1986

This Technical Note formerly described a bug of SANE on the Apple IIGS which caused it to jump through location \$00/0018 instead of through the HALT vector in the SANE direct page.

The bug which caused SANE on the Apple IIGS to jump through location \$00/0018 instead of through the HALT vector in the SANE direct page was fixed in the Apple IIGS ROM 2.0. You should not have to write a special case to handle this bug since it is reasonable to expect users to have the updated ROM which is offered as a free upgrade from Apple.

Apple IIGS

#8: Elems Functions in IIGS SANE

Revised by: Matt Deatherage

November 1988

Written by: Guillermo Ortiz

December 1986

This Technical Note discusses a problem which existed with the Elems functions in the IIGS SANE Tool Set 1.0. Current IIGS System Disks contain a patch which corrects this problem.

Calls to any of the Elems functions in version 1.0 of the IIGS SANE Tool Set may return an invalid result unless you are evaluating data which resides in bank \$00 due to a problem with the Elems parameter passing mechanism. These results are random because when SANE checks the validity of its input, it uses values that have no relations to the actual ones, and once it completes the validation, it uses the real operands.

All System Disks released on or after December 1, 1986 include a RAM patch which fixes the Elems parameter passing mechanism; therefore, you should not have to write a special case to handle this problem if you are shipping your application with the most recent Apple IIGS System Disk. You should contact Apple Software Licensing at Apple Computer, Inc.; 20525 Mariani Avenue, M/S 38-I; Cupertino, CA 95014 or (408) 974-4667 to obtain the most recent version of the Apple IIGS System Disk.

Further Reference

- o Apple Numerics Manual

Apple IIGS

#9: IIGS Sound Expansion Connector:
Analog Input/Output Impedances

Revised by: Jim Merritt & Jim Mensch
Written by: Jim Merritt

November 1988
December 1986

This Technical Note discusses the impedances of the analog signal pins on the IIGS sound expansion connector since an interface to this connector must take the impedance of the pins into account to function properly.

The analog output impedance of pin 3 depends upon the characteristics of the 5503 sound synthesis chip in any particular IIGS machine. Across systems, this impedance may range from 4.5 K ohms to 9 K ohms.

Pin 1, the A/D input, presents a dynamic load to the source, drawing at 10 K ohms for approximately 500 ns during every sample period. It is reasonable, however, to treat the input pin as if it presents a continuous load of 10 K ohms without compromising the interface or the fidelity of the input sample.

Consult the Apple IIGS Hardware Reference for further technical information about the Ensoniq 5503 sound synthesis chip used in the IIGS.

Further Reference

- o Apple IIGS Hardware Reference

Apple IIGS
#10: InvalRgn Twist

Revised by: Steven Glass
Written by: Guillermo Ortiz

November 1988
April 1987

InvalRgn(RgnHandle) accumulates the region to which RgnHandle points into the update region of the current window's port; in the process, it makes the region global, thus causing problems if later calls expect the region to still be local.

The region you pass to InvalRgn is local to the window to which it is related; however, InvalRgn returns the region in global coordinates. To preserve the original region for your use after the call to InvalRgn, you should duplicate it and use the copy to make the call then dispose of the copy when InvalRgn returns. The following example demonstrates the process:

```
void MyInvalReg(RegHandle)

handle RegHandle;
{
handle AuxHandle;

AuxHandle = NewRgn();          /* create room */
CopyRgn(RegHandle,AuxHandle); /* make a copy */
InvalRgn(AuxHandle);          /* do it with the copy */
DisposeRgn(AuxHandle);       /* now get rid of it! */
}
```

Further Reference

- o Apple IIGS Toolbox Reference, Volume 2

Apple IIGS

#11: Ensoniq DOC Swap-Mode Anomaly

Revised by: Jim Mensch
Written by: Jim Merritt

November 1988
April 1987

Under certain conditions, the IIGS Ensoniq Digital Oscillator Chip (DOC) inserts a spurious zero-crossing byte into the output sample stream. The output sample waveform may mask the anomaly, but if it does not, the user may hear intermittent clicks or even a more pervasive "static." This Technical Note discusses the situations in which the DOC produces this spurious zero crossing, as well as strategies to avoid or mask this undesirable behavior.

Background

The Ensoniq DOC in the Apple IIGS is actually a microprocessor dedicated to producing sound. Like a time-sharing computer, the DOC continually scans through its array of sound oscillators, proceeding from lower-numbered oscillators to higher-numbered ones, and updates the signal output level of each active one to match that indicated by the oscillator's current sample byte.

An oscillator can operate in any one of several functional modes, as described in the Apple IIGS Hardware Reference. Here, however, we are concerned only with swap mode, where two consecutive oscillators are considered as a single generator. The low-numbered oscillator in the pair is always even. For example, the pairs of oscillators 0 & 1, 2 & 3, ... , 12 & 13, and 14 & 15 constitute generators. The IIGS Sound Tool Set - the FFStartSound call in particular - configures the oscillators it uses to operate in swap mode. In swap mode, the even-numbered oscillator plays its waveform first, halts its own playback, then starts its partner which also plays its waveform, halts its own playback upon exhausting its waveform, and restarts the even-numbered oscillator. At any time between the start of any particular FFStartSound call and the time the oscillator finishes playing a wave, the Sound Tool Set interrupt handler may be busy transferring waveform information from the IIGS main RAM to the dormant oscillator's buffer in DOC RAM. Since one oscillator is producing sound while the Sound Tool Set interrupt handler is transferring waveform information to the other oscillator, you can use a generator pair to produce continuous sound of arbitrary length, and you are limited only by the amount of memory you can devote to the waveform in the main RAM.

Each oscillator draws its output samples from a dedicated buffer in DOC RAM, the size and location of which are specified by parameters to the FFStartSound call. The maximum size for an oscillator buffer is 32K, but since buffers may neither coincide nor overlap, the practical maximum may be lower when more than one generator is active. For instance, if four generators (eight paired oscillators) are active simultaneously, the maximum buffer size is 8K, since eight non-overlapping buffers of 8K each would occupy the entire 64K available in the DOC RAM.

The Problem

Whenever a swap occurs from a higher-numbered oscillator to a lower-numbered

one, the output signal from the corresponding generator temporarily falls to the zero-crossing level (silence); this anomaly does not occur during swaps from lower-numbered oscillators to higher-numbered ones. The spurious level change lasts no longer than a single sample period, at which time the interrupted waveform resumes. However, even this tiny glitch in the output can be audible as a pop or click; the further away the waveform is from the zero crossing when the swap interrupts it, the louder the ear will perceive the pop or click. When high-to-low swaps occur with great frequency, the pops and clicks happen so often that they are perceived as gentle, but pervasive, static.

Several Workarounds

There is no ideal solution to the problem of signal interruption in swap mode. This problem is an anomaly of the DOC design, which may or may not be addressed in later versions of the chip. However, we have found three general strategies for mitigating the audible damage to the output waveform caused by the chip's undesirable behavior.

Minimize Oscillator Swaps per Unit Time

The more often swaps from high-numbered oscillators to low-numbered ones occur, the more obtrusive the brief signal interruptions will seem. To minimize the interruptions, you must make the oscillators play for a longer period of time before swapping to their partners. This means that they must play at slower output sample rates, use larger buffers in DOC RAM, or use the two in tandem. Commensurate with the number of active generators you wish to use and the level of output signal fidelity that you desire, always specify the largest DOC buffer size and the lowest output sample rate that you possibly can. Remember that a large number of active generators implies a very small maximum buffer size for any particular oscillator, so you should always try to minimize the number of generators that are active at any one time. As a rough benchmark, the clicks of signal interruption begin to blend into highly audible static when you specify buffers smaller than 8K for use at the maximum-fidelity output sample rate of about 26 kHz. (Note: The DOC supports greater sample rates, but these rates are limited by the output filtering on the IIGS which permits no greater signal fidelity than that possible using the 26 kHz rate.) Our figures suggest that output fidelity must suffer, or signal noise must increase, when more than four generators (eight oscillators in swap mode) are operating simultaneously.

Avoid Silent or Quiet Passages

The signal content of your waveform can hide the additional noise caused by the "swap-mode anomaly." The more complex and louder a waveform, the less your ear will perceive the brief interruption that occurs whenever a higher-numbered oscillator swaps to a lower-numbered one; pop and rock music is far less susceptible to this problem than classical, folk, or jazz pieces, which typically include many quiet passages. In addition, a signal that naturally contains a large amount of "pink noise," such as recordings of rainstorms or the surf at the beach, can mask the anomalous noise altogether.

Arrange for Swaps to Occur at or Near Zero Crossings

If the high-to-low swap occurs at a time when the normal output signal level sits at or near the zero crossing, the swap will cause little or no audible damage to the waveform. When reproducing arbitrary sampled sound, it is almost impossible to insure that the output signal level is near the zero crossing. However, when constructing long waveforms for playback, you may be able to sidestep the chip's anomalous behavior by ensuring that the waveform values lie at or near zero at the end of every waveform segment, where a

waveform segment spans twice the length of one oscillator buffer. For example, if you specify a buffer size of 4K, make sure that your constructed waveform crosses the baseline after every 8,192 samples, and for 16K buffers, make sure that the waveform makes a zero crossing after every 32K.

The length of the waveform segment should be twice the buffer length only if you are going to reproduce the waveform exactly once per FFStartSound call. It may be necessary to shorten the length of the waveform segment to exactly the specified DOC buffer length if you use the nextwave_start parameter in the FFStartSound parameter block to invoke automatic looping of the waveform. In other words, you may need to arrange for twice as many zero crossings in your constructed waveform in the looping case as you would under normal circumstances since subsequent repetitions of the waveform during the single FFStartSound call may begin with either the even or odd oscillator, depending upon which member of the pair was active when the previous repetition ended. If the playback of a waveform starts with the odd oscillator, then the odd-to-even swaps will occur at different points in the waveform than they would when the playback starts with the even oscillator.

Also note that the use of larger buffers causes a progressively longer disabling of interrupts while the Sound Tool Set moves the waveform into the DOC RAM.

Further Reference

- o Apple IIGS Toolbox Reference, Volume 2
- o Apple IIGS Hardware Reference

Apple IIgs
#12: Tool Set Interdependencies

Revised by: Matt Deatherage & Dave Lyons
Written by: Jim Merritt

May 1992
April 1987

This Technical Note lists all known interdependencies between system tool sets on the Apple IIgs.

CHANGES SINCE JANUARY 1990: Added new and changed dependencies for System Software 6.0.

A tool set is dependent upon another if you must start the latter before starting the former. You should start tool sets in the order listed below. Names marked with an asterisk (*) indicate a recommendation to start the corresponding tool set, but the order is not required for operation of the dependent tool. Apple recommends using StartUpTools to start up all the tool sets your application needs. See the Apple IIgs Toolbox Reference, Volume 3 for more details.

TOOL SET INTERDEPENDENCIES

Tool Locator Tool #1 (\$01)
No dependencies. Always start this tool set before any others.

Memory Manager Tool #2 (\$02)
Tool Locator (#1)

Miscellaneous Tools Tool #3 (\$03)
Tool Locator (#1)
Memory Manager (#2)

QuickDraw II Tool #4 (\$04)
Tool Locator (#1)
Memory Manager (#2)
Miscellaneous Tools (#3)

Desk Manager Tool #5 (\$05)
Tool Locator (#1)
Memory Manager (#2)
Miscellaneous Tools (#3)
QuickDraw II (#4)
Event Manager (#6)
Window Manager (#14)
Control Manager (#16)
Menu Manager (#15)
Line Edit (#20)
Dialog Manager (#21)
Scrap Manager (#22)

Event Manager Tool #6 (\$06)
Tool Locator (#1)
Memory Manager (#2)
Miscellaneous Tools (#3)

Scheduler Tool #7 (\$07)

Tool Locator	(#1)	
Memory Manager	(#2)	
Miscellaneous Tools	(#3)	
Sound Tools Set		Tool #8 (\$08)
Tool Locator	(#1)	
Memory Manager	(#2)	
Miscellaneous Tools	(#3)	
Apple Desktop Bus (ADB)		Tool #9 (\$09)
Tool Locator	(#1)	
SANE (Standard Apple Numeric Environment)		Tool #10 (\$0A)
Tool Locator	(#1)	
Memory Manager	(#2)	
Integer Math Tools		Tool #11 (\$0B)
Tool Locator	(#1)	
Text Tools		Tool #12 (\$0C)
Tool Locator	(#1)	
Window Manager		Tool #14 (\$0E)
Tool Locator	(#1)	
Memory Manager	(#2)	
Miscellaneous Tools	(#3)	
QuickDraw II	(#4)	
Event Manager	(#6)	
* QuickDraw Aux	(#18)	Required in 6.0 and later, and the window manager loads and starts it for you.
Control Manager	(#16)	
Menu Manager	(#15)	
* Line Edit	(#20)	For AlertWindow call only
* Font Manager	(#27)	For AlertWindow call only
* Resource Manager	(#30)	For using resources in Window Manager calls.
Menu Manager		Tool #15 (\$0F)
Tool Locator	(#1)	
Memory Manager	(#2)	
Miscellaneous Tools	(#3)	
QuickDraw II	(#4)	
Event Manager	(#6)	
Window Manager	(#14)	
Control Manager	(#16)	
* Resource Manager	(#30)	For using resources in Menu Manager calls.
Control Manager		Tool #16 (\$10)
Tool Locator	(#1)	
Memory Manager	(#2)	
Miscellaneous Tools	(#3)	
QuickDraw II	(#4)	
Event Manager	(#6)	
Window Manager	(#14)	
Menu Manager	(#15)	
* QuickDraw Auxiliary	(#18)	For statText controls.
* Line Edit	(#20)	For editLine controls.
* Font Manager	(#27)	For statText controls.

- * List Manager (#28) For list controls.
- * Resource Manager (#30) For using resources in Control Manager calls.
- * Text Edit (#34) For editText controls.

NOTE: You should consider the Window, Control, and Menu Managers as one unit and start them in the given order.

System Loader Tool #17 (\$11)

- Tool Locator (#1)
- Memory Manager (#2)
- Miscellaneous Tools (#3)

QuickDraw Auxiliary Routines Tool #18 (\$12)

- Tool Locator (#1)
- Memory Manager (#2)
- Miscellaneous Tools (#3)
- QuickDraw II (#4)
- * Font Manager (#27)

NOTE : QuickDraw Auxiliary uses the Font Manager in the picture drawing routines. For proper operation, you should start the Font Manager before using the QuickDraw Auxiliary picture routines; however, the picture routines do not fail if the Font Manager is not present.

Print Manager Tool #19 (\$13)

- Tool Locator (#1)
- Memory Manager (#2)
- Miscellaneous Tools (#3)
- QuickDraw II (#4)
- QuickDraw Auxiliary (#18)
- Event Manager (#6)
- Window Manager (#14)
- Control Manager (#16)
- Menu Manager (#15)
- Line Edit (#20)
- Dialog Manager (#21)
- List Manager (#28)
- Font Manager (#27)

Line Edit Tool #20 (\$14)

- Tool Locator (#1)
- Memory Manager (#2)
- Miscellaneous Tools (#3)
- QuickDraw II (#4)
- Event Manager (#6)
- * QuickDraw Auxiliary (#18) For Text2 items; see below.
- Scrap Manager (#22)
- * Font Manager (#27) For Text2 items; see below.

Dialog Manager Tool #21 (\$15)

- Tool Locator (#1)
- Memory Manager (#2)
- Miscellaneous Tools (#3)
- QuickDraw II (#4)
- Event Manager (#6)
- Window Manager (#14)
- Control Manager (#16)
- Menu Manager (#15)
- * QuickDraw Auxiliary (#18) For Text2 items; see below.

Line Edit (#20)
* Font Manager (#27) For Text2 items; see below.

NOTE : Line Edit, the Dialog Manager, and the Control Manager require the presence of the Font Manager and QuickDraw Auxiliary if you use LERectBox2, statText controls, or LongStatText2 items which require any font styling (e.g., outline, boldface, etc.).

Scrap Manager Tool #22 (\$16)
Tool Locator (#1)
Memory Manager (#2)

Standard File Operations Tool #23 (\$17)

Tool Locator (#1)
Memory Manager (#2)
Miscellaneous Tools (#3)
QuickDraw II (#4)
Event Manager (#6)
Window Manager (#14)
Control Manager (#16)
Menu Manager (#15)
* QuickDraw Auxiliary (#18)

Required in 6.0 and later,
and the Window Manager loads
and starts it for you.

Line Edit (#20)
Dialog Manager (#21)
* List Manager (#28)
* Resource Manager (#30)

For using resources in
Standard File Operations
calls.

NOTE : Standard File 3.0 and later use the List Manager for displaying a list of file names. Although Standard File functions properly if the application has not started the List Manager, it saves time if the application does so.

Note Synthesizer Tool #25 (\$19)
Tool Locator (#1)
Memory Manager (#2)
Sound Tools (#8)

Note Sequencer Tool #26 (\$1A)
Tool Locator (#1)
Memory Manager (#2)
Sound Tools (#8)
Note Synthesizer (#25)

Note : The Note Sequencer automatically handles the start and shutdown of the Free-Form Sound Tools (#8) and the Note Synthesizer (#25), so programs that use the Note Sequencer must not execute start or shutdown calls for those tools. Automatic start does not imply automatic loading. If you plan to use the Note Sequencer, you must still load the Free-Form Sound Tool and the Synthesizer Tool explicitly through calls to the Tool Locator routines LoadTools or LoadOneTool or by calling the System Loader and Tool Locator directly in appropriate cases.

Font Manager Tool #27 (\$1B)

Tool Locator	(#1)	
Memory Manager	(#2)	
* Miscellaneous Tools	(#3)	For ChooseFont call only.
QuickDraw II	(#4)	
* Integer Math Tools	(#11)	For ChooseFont call only.
* Window Manager	(#14)	For ChooseFont call only.
* Control Manager	(#16)	For ChooseFont call only.
* Menu Manager	(#15)	For FixFontMenu call only.
* List Manager	(#28)	For FixFontMenu and ChooseFont calls.
* Line Edit	(#20)	For ChooseFont call only.
* Dialog Manager	(#21)	For ChooseFont call only.
List Manager		Tool #28 (\$1C)
Tool Locator	(#1)	
Memory Manager	(#2)	
Miscellaneous Tools	(#3)	
QuickDraw II	(#4)	
Event Manager	(#6)	
Window Manager	(#14)	
Control Manager	(#16)	
Menu Manager	(#15)	
Audio Compression and Expansion (ACE)		Tool #29 (\$1D)
Tool Locator	(#1)	
Memory Manager	(#2)	
Resource Manager		Tool #30 (\$1E)
Tool Locator	(#1)	
Memory Manager	(#2)	
MIDI Tools		Tool #32 (\$20)
Tool Locator	(#1)	
Memory Manager	(#2)	
Miscellaneous Tools	(#3)	
Sound Manager	(#8)	
* Note Synthesizer	(#25)	
NOTE : The MIDI Tools require the Note Synthesizer if you intend to use the MIDI clock feature. If you are not using the MIDI clock, the Note Synthesizer is not required.		
Text Edit		Tool #34 (\$22)
Tool Locator	(#1)	
Memory Manager	(#2)	
Miscellaneous Tools	(#3)	
QuickDraw II	(#4)	
Event Manager	(#6)	
Window Manager	(#14)	
Menu Manager	(#15)	
Control Manager	(#16)	
QuickDraw Auxiliary	(#18)	
Scrap Manager	(#22)	
Font Manager	(#27)	
* Resource Manager	(#30)	For using resources in Text Edit calls.
MIDI Synth		Tool #35 (\$23)
Tool Locator	(#1)	

Memory Manager (#2)
 Miscellaneous Tools (#3)
 Sound Tools (#8)

Media Control Tool Tool #38 (\$26)
 Tool Locator (#1)
 Memory Manager (#2)
 Miscellaneous Tools (#3)
 Integer Math (#11)
 Resource Manager (#30)

Recommended Start Order

A close look at the preceding information will reveal apparent "circular dependencies" between various tool sets (i.e., two or more tool sets may depend upon each other). To resolve the issue of which tool set to start first in such a situation, here is a list of the most commonly used tool sets, given in the order in which an application should start them. You may start those tools which are indented at a specific level at that time or any time thereafter.

Tool Locator	(#1)		
	ADB Tools	(#9)	
	Integer Math Tools	(#11)	
	Text Tools	(#12)	
Memory Manager	(#2)		
	SANE	(#10)	
	ACE	(#29)	
Resource Manager	(#30)		
Miscellaneous Tools	(#3)		
	Scheduler	(#7)	
	System Loader	(#17)	LoaderStartup does nothing.
	Media Control	(#38)	
QuickDraw II	(#4)		
	QuickDraw II Auxiliary	(#18)	
Event Manager	(#6)		
Window Manager	(#14)		
Control Manager	(#16)		
Menu Manager	(#15)		
Line Edit	(#20)		
Dialog Manager	(#21)		
either			
	Sound Tools then	(#8)	
	Note Synthesizer	(#25)	
or			
	Note Sequencer	(#26)	
	MIDI Tools	(#32)	
	MIDI Synth	(#35)	
Standard File	(#23)		
Scrap Manager	(#22)		
List Manager	(#28)		
Font Manager	(#27)		
Print Manager	(#19)		
Text Edit	(#34)		
Desk Manager	(#5)		

NOTE : Although you may start the sound-related tools any time after the Miscellaneous Tools, we recommend you start them after most of the Desktop-related tools. We also recommend you start the Desk Manager last and shut it

down first.

Further Reference

- o Apple IIgs Toolbox Reference

Apple IIGS
#13: ROM 1.0 Modem Firmware Bug

Revised by: Matt Deatherage
Written by: Mike Askins

November 1988
April 1986

This Technical Note formerly discussed a bug involving buffering and serial port setting commands in the modem firmware in ROM 1.0.

Apple IIGS ROM 2.0 fixes a bug involving buffering and serial port setting commands in the modem firmware. You should not have to write a special case to handle this bug since it is reasonable to expect users to have the updated ROM which is offered as a free upgrade from Apple.

Apple IIgs
#14: Standard File Screwiness

Revised by: Dave Lyons
Written by: Guillermo Ortiz, Matt Deatherage, & Dave Lyons

May 1992
June 1987

This Technical Note describes known anomalies in Standard File.

CHANGES SINCE DECEMBER 1991: Updated for System 6.0. Problems with the infinite loop and SFMultiGet2 reply record are fixed.

PREFIX CHECK IS CASE SENSITIVE

When you advance to the next volume using Command-Tab (or just Tab, before 6.0), Standard File checks your prefix against the name of the volume now in the same device you were just using, to see if you switched disks (this is possible on a 5.25 drive, for example). If the name doesn't match, you stay at the same device.

Unfortunately, the comparison in 6.0 and earlier is case sensitive. If you have a volume called "MyDisk" and prefix zero is set to ":MYDISK", advancing to the next volume doesn't get you anywhere the first time (but the prefix changes from ":MYDISK" to ":MyDisk").

The following two problems are fixed in System 6.0:

INFINITE LOOP WITH EMPTY PREFIXES

In System Software versions 5.0 through 5.0.4, all Standard File dialogs can hang if both prefixes 0 and 8 are empty (GS/OS uses prefix 8 to expand partial pathnames if prefix 0 is empty).

If this affects your software, use GetPrefix to check for empty prefixes before calling Standard File. If 0 and 8 are both empty, set prefix 0 to "':" (or any other convenient pathname).

SFMultiGet2 (AND SFPMultiGet2) REPLY RECORD

SFMultiGet2 and SFPMultiGet2 in System 5.0.4 and earlier accidentally validate the multi-file reply record as if it were a regular new-style reply record (as for SFGetFile2, for example). The validation is a check that the words at offsets \$08 and \$0E in the record are not \$0002 (these are nameRefDesc and pathRefDesc in a new-style reply record).

To ensure that Standard File does not erroneously reject your multi-file reply record (and return error \$1704), you may include ten bytes of \$00 following the six-byte record.

Further Reference

-
- o Apple IIgs Toolbox Reference, Volumes 2 & 3

Apple IIGS

#15: InstallFont and Big Fonts

Revised by: Eric Soldan & Matt Deatherage

July 1989

Written by: Guillermo Ortiz

June 1987

When the Font Manager executes InstallFont, it may try to scale the selected font if bit 15 of the ScaleWord is clear; a font larger than 32K causes this call to fail.

Changes since November 1988: Noted System Software 5.0 enhancements.

The Font Manager cannot scale a font which is larger than 32K, so InstallFont will fail if scaling is required and the desired font exceeds this limit. If the call fails for this reason, it will report an FMScaleSizeErr (\$1B0C) error.

This is not the same situation as when there is not enough memory available to hold a newly scaled font. The situation will generate Memory Manager errors.

System Software 5.0 can scale fonts to be larger than 32K, so there is no longer the limit imposed by System Disk 4.0 and earlier. In addition, System Software 5.0 can handle font sizes up to 255 points, if memory is available. Note that this is a different situation than trying to scale a font which was originally larger than 32K, but both work under 5.0.

Apple IIGS
#16: Notes on Background Printing

Revised by: Mike Askins
Written by: Mike Askins

November 1988
June 1987

This Technical Note attempts to pinpoint some of the common problems people encounter when using background printing as available through the serial firmware.

Calling Sequence

Init call	Starts the serial firmware
SetOutBuff	Specifies a buffer to place data to be printed
	Places data in buffer (amount < buffer size)
SendQueue	Starts the background printing process

Correctly Making the SendQueue Call

The Apple IIGS Firmware Reference incorrectly documents the parameters you pass to SendQueue. The correct specification of the recharge address does not correspond to the standard method of passing a full 32-bit address. Set the parameters as follows:

SendQueue
Launches background printing.

CmdList	DFB \$04	;Parameter Count
	DFB\$18	;Command Code
	DW \$00	;Result Code (output)
	DW DataLength	
	DFB RechargeAddress (bank)	
	DFB RechargeAddress (high)	
	DFB RechargeAddress (low)	
	DFB \$00	

Using the Default Buffer

You can use the area which the firmware reserves for transparent buffering to place data for background printing. This is advantageous since the firmware calls the Memory Manager to allocate space for the buffer (you must allocate the space from the Memory Manager if you use the SetOutBuff call to set up a buffer).

To use the serial firmware's buffer, you must first enable buffering by initializing the port with PINIT and sending it the string "^IBE" with PWRITE. Once you enable buffering, call GetOutBuff to find the size and location of the buffer, then place your data (buffer size - 1) in the buffer and call SendQueue.

Data Size

Make sure that the amount of data you place in the buffer is at least one byte less than the size of the buffer since the firmware uses one byte of the buffer for bookkeeping purposes; if you place too much data in the buffer, it will continually print the buffer's contents and never call your recharge routine.

The Recharge Routine

You should treat the recharge routine as an interrupt handler and execute it at interrupt time. Interrupts are disabled at this time, and it is illegal to enable them within the recharge routine. Like all interrupt handlers, the recharge routine should take care of its business as quickly as possible then exit; any excessive delays cause interrupt dependent processes (e.g., AppleTalk) to fail. You should also remember that most of the system code is non-reentrant; you should use the Scheduler when calling system code which may have been running when the serial interrupt that invoked the recharge routine occurred.

The serial firmware is not generally reentrant and does not interact with the Scheduler. If you want to make serial firmware calls (through \$C1xx, \$C2xx) from your recharge routine, you must preserve MSLOT (the byte at \$0007F8) across those calls. Be aware that any non-recharge code must not make calls to the serial firmware that will disrupt the background printing process; sending the string "^BD" (disable buffering command), for example, is guaranteed to confuse a running background printing process.

Further Reference

- o Apple IIGS Firmware Reference

Apple IIGS

#17: Application Memory Management and the MMStartUp User ID

Revised by: Steven Glass & Rich Williams

November 1988

Written by: Jim Merritt

June 1987

This Technical Note describes a technique which permits an application to dispose of any memory it has used with a single Memory Manager call without clobbering other system components or itself.

Ground Rules for Application Memory Usage

Apple IIGS programs must be responsible for allocating and disposing of any memory they use, over and above that which the operating system itself gives them. In general, no IIGS program should use any memory except that which the Memory Manager has explicitly granted to it. A program may request additional memory for its own use at any time with one or more calls to the NewHandle routine. At program termination, the application is responsible for explicitly disposing of any memory that it explicitly acquired, and if it fails to do so, it could leave the IIGS memory management system in a corrupted state.

You may dispose of memory on a handle-by-handle basis, or you may dispose of it en masse by calling DisposeAll, but you should never use DisposeAll with the user ID that the MMStartUp routine provides. This user ID is the "master user ID" for the application, and it tags the memory space which the operating system reserves for the program's code and static data at load time. Calling DisposeAll with this user ID results in immediate deallocation of the memory in which the calling program resides; therefore, an application which allocates dynamic data space using only the user ID that MMStartUp gives it should not use DisposeAll to deallocate that space, but rather use DisposeHandle to deallocate it handle by handle.

Cleaning Up With DisposeAll

It is possible, however, for a program to use a different, unique user ID when allocating its own RAM, then pass that user ID to DisposeAll when it terminates to deallocate all of its private memory at once without endangering itself or other parts of the IIGS system. With this technique, the question is how best to acquire a new user ID? One method to acquire a new user ID is to request a completely new one of the appropriate type from the User ID Manager in the Miscellaneous Tools. In this case, when the application terminates, it must not only deallocate the memory it used, but also the additional user ID which it requested from the User ID Manager.

Actually, it is not necessary for a program to acquire a completely new user ID to use DisposeAll without clobbering itself. Instead, the application may modify the auxID field of the master user ID which MMStartUp assigns to create a unique user ID for allocating its own memory. The 16-bit user ID contains the auxID field in bits \$8 - \$B. The value of this field, which may range from \$0 to \$F, is always zero in the application's master user ID, but you can fill it with any non-zero value to create up to 15 new and distinct user IDs, each of which you can pass to NewHandle to allocate memory and to DisposeAll

to deallocate memory without endangering the memory tagged by the master user ID. The following assembly code fragment illustrates this technique:

```
; assumes full native mode
    pushword #0                ; room for user ID
    _MMStartUp
    pla                        ; master user ID
    sta MasterID
    ora #$0100                ; auxID:= 1

;   (COULD HAVE BEEN ANYTHING FROM $1 to $F)

    sta MyID                  ; use this to allocate private memory
    ...
    etc.
    ...

; ready to exit program
    pushword MyID
    _DisposeAll               ; dumps only my own RAM

; now do any remaining processing related to termination
```

You do not need to explicitly deallocate any user ID that you derive by changing the auxID field of a valid master user ID. When the system (usually the one to deallocate the master) deallocates the master user ID, it also deallocates its derivatives.

One Word of Caution

Several of the Memory Manager's "All" calls (e.g., DisposeAll) treat a zeroed auxID field as a wildcard which matches any value that the field may contain, thus if you call DisposeAll with the application's master user ID (where the auxID field is zero), the Memory Manager will not only deallocate all memory belonging to the master user ID, but also all handles and memory segments that are associated with user IDs which are derived from that master. The operating system's QUIT mechanism typically executes such a call when cleaning up after a normal (i.e., non-restartable) application to keep the memory management system from clogging. This action is purely a defensive measure, and well-behaved applications - particularly restartable ones - should dispose of their own memory and never rely upon the operating system to clean up after them.

Further Reference

- o Apple IIGS Toolbox Reference, Volume 1

Apple IIgs

#18: Do-It-Yourself SCC Access

Revised by: Jim Luther

July 1990

Written by: Jim Luther, Mike Askins, Matt Deatherage & Jim Mensch June 1987

This Technical Note describes how to install and remove an interrupt handler routine for the Z8530 Serial Communications Controller (SCC) on the Apple IIgs without breaking other parts of the system. This Note includes many suggestions that, if unheeded, could come back to haunt you in the form of bug fixes to your program.

Changes since March 1990: Added a method for finding which serial port AppleTalk is using under GS/OS.

Free Serial Routines Inside

The Z8530 SCC has 2 serial channels, supports several synchronous and asynchronous data communications protocols, and has 9 read registers and 16 write registers per channel. (Compare this to the 5 registers of the 6551 Asynchronous Communications Interface Adapter.) To program the SCC correctly, you must understand five things: the SCC, the Apple IIgs hardware environment in which the SCC lives, the Apple IIgs interrupt handler firmware, the interrupt support provided by the operating system, and the data communication protocol you want to use. If you don't understand all of these components, stick to the serial firmware.

The Apple IIgs serial firmware is a robust environment for almost every asynchronous serial programming application. If you want to handle all SCC operations and SCC interrupts on the IIgs without using the serial firmware, then you must really know the firmware won't do the job for you or you wouldn't be going to a lot of trouble to recreate the services the firmware routines already provide.

Don't Eat Your Serial with Your Mouth Open

Your mother has rules and so does Apple. On many systems, your application may be sharing the SCC chip with System Software such as AppleTalk or the serial firmware. If you want to access the SCC chip directly without breaking the system (or the system breaking you), then follow these simple rules.

Rule #1: Before using a serial port, make sure AppleTalk is not already using it.

If AppleTalk is active, it uses one of the serial ports. The user selects which serial port AppleTalk uses with the Control Panel. Before using one of the serial ports, you should always check to make sure AppleTalk is not using that port. If AppleTalk is using the serial port your application wants to use, tough luck; tell the user about it, but don't even think about using that port.

Under ProDOS 8, use the method shown in the following sample code to determine if AppleTalk is using a serial port:

```
;  
; This routine checks to see which serial port, if any, AppleTalk is using.
```

```

; The routine sets a flag byte, ApTalkPort, and the accumulator to indicate
; which port (if any) AppleTalk is using.
;   $00 = AppleTalk is not using a serial port
;   $01 = AppleTalk is using serial port 1 (printer port)
;   $02 = AppleTalk is using serial port 2 (modem port)
; Note: This method should be used under ProDOS 8 only. Under GS/OS, use the
;       .AppleTalk driver's GetPort DStatus subcall.
;
; Enter routine in emulation mode
;
                longa off
                longi off
                mcopy 2/AInclude/M16.MiscTool

WhichPort      start

IDROUTINE      equ $FE1F                returns system ID information

                stz ApTalkPort          default to not AppleTalk

                jsr IDROUTINE           call to the system ID routine
                cpy #$03
                bcs NewIIGS

OldIIGS        anop                    this is a pre-ROM 03 IIGS
                clc                    to native mode
                xce
                rep #$30                16 bit m and x
                longa on
                longi on

                pea $0000               space for result
                pea $0021               Slot 1 setting
                _ReadBParam             read battery RAM parameter
;                                     (2 byte result left on stack)

                pea $0000               space for result
                pea $0027               Slot 7 setting
                _ReadBParam             read battery RAM parameter
                pla                     get slot 7 setting (2 bytes)

                sec                    emulation mode
                xce
                longa off
                longi off

                beq FindYourCard        AppleTalk is active
                pla                     remove slot 1 setting LSB (1 byte)
                bra OldExit

FindYourCard   inc ApTalkPort          default to port 1
                pla                     is slot 1 "your card"? (1 byte)
                beq ItsPort2            no, must be port 2
                bra OldExit

ItsPort2      inc ApTalkPort          port 2 is AppleTalk

OldExit       pla                     remove slot 1 setting MSB (1 byte)
                lda ApTalkPort
                rts                    return to caller

NewIIGS       anop                    ROM 03 or greater IIGS

```

```

        clc                to native mode
        xce
        rep #$30          16 bit m and x
        longa on
        longi on

        pea $0000         space for result
        pea $000C         port 2 type
        _ReadBParam      read battery RAM parameter
                          (2 byte result left on stack)
;

        pea $0000         space for result
        pea $0000         port 1 type
        _ReadBParam      read battery RAM parameter
        pla              get port 1 setting (2 bytes)

        sec              emulation mode
        xce
        longa off
        longi off

        cmp #$02         is port 1 AppleTalk?
        bne TryPort2     no
        inc ApTalkPort   yes
        pla              then remove port 2 setting LSB (1 byte)
        bra NewExit      and exit

TryPort2        pla      get port 2 setting LSB (1 byte)
                cmp #$02 is port 2 AppleTalk?
                bne NewExit no
                lda #$02 yes
                sta ApTalkPort

NewExit        pla      remove port 2 setting MSB (1 byte)
                lda ApTalkPort
                rts      return to caller

ApTalkPort    entry
                ds 1    will be 0, 1, or 2
                end

```

Under GS/OS, use the method shown in the following sample code to determine if AppleTalk is using a serial port:

```

;
; This routine checks to see which serial port, if any, AppleTalk is using.
; The routine sets a flag byte, ApTalkPort, and the accumulator to indicate
; which port (if any) AppleTalk is using.
; $0000 = AppleTalk is not using a serial port
; $0001 = AppleTalk is using serial port 1 (printer port)
; $0002 = AppleTalk is using serial port 2 (modem port)
; Note: This method should be used under GS/OS only.
;
; Enter routine in native 16 bit mode
;
        longa on
        longi on
        mcopy 2/AInclude/M16.GSOS

CheckPort    Start

GetPort      equ $8001          The .AppleTalk DStatus subcall to get

```



```

DSdevNum          ds 2          devNum
                  dc i2'GetPort'  statusCode = GetPort
                  dc a4'GetPortSList'  statusList = GetPortSList
                  dc i4'2'        requestCount = 2
                  ds 4          transferCount

GetPortSList      anop          the GetPort subcall's statusList
portNum           ds 2          $0001 = AppleTalk is using port 1 (printer port)
;                                                         $0002 = AppleTalk is using port 2
(modem port)
                  dc i2'0'

                  end

```

Rule #2: Don't use the SCC Interrupt Handler Vector.

Contrary to what you may have read in a previous version of this Note, you cannot reliably attach your SCC interrupt handler to the SCC Interrupt Handler Vector (vector reference number \$0009). The Apple IIgs serial firmware owns the SCC Interrupt Handler Vector (or at least it thinks it does). Anytime the serial firmware is used, there is a chance that the serial firmware can grab the SCC Interrupt Handler Vector for its use. CDAs and NDAs that print, the Print Manager tool set, the Text tool set, and the generated GS/OS character drivers associated with the serial ports are examples of code that can and do use the serial firmware.

The only safe place to connect into the interrupt chain is through the operating system. The ProDOS 8 and GS/OS ProDOS 16 call, `ALLOC_INTERRUPT` is the correct place to attach your interrupt handler. The GS/OS `BindInt` call cannot be used to attach your interrupt handler to the SCC Interrupt Handler Vector (VRN \$0009) for the same reason that you cannot use the SCC Interrupt Handler Vector directly.

Rule #3: Be very, very careful with SCC Write Register 9 (WR9).

The Z8530 SCC has four registers which are shared by both channels (ports). Of those four, only two are commonly used in the Apple IIgs, RR3 and WR9. RR3, which only exists in channel A, lets you check the interrupt pending bits for both SCC channels. WR9 is the Master Interrupt Control register for both SCC channels and contains the Reset command bits.

You must never reset the channel AppleTalk is using (resetting the channel AppleTalk is using kills AppleTalk). This means you should never perform a Force Hardware Reset command (11xxxxxx to WR9) even though the Z8530 Serial Communications Controller Technical Manual tells you to in the SCC initialization procedure. A hardware reset is performed at system startup, so you shouldn't need to perform a channel reset, even to the channel you are using.

The interrupt control bits (bits D5 - D0) in WR9 should not be modified (an exception is when you are installing your own SCC interrupt handler). AppleTalk expects the interrupt control bits to always be 001010. If you find the need to perform a channel reset on your channel, remember that the interrupt control bits are programmed at the same time as a channel reset.

Hints for the Serial Adventure

Next are a few hints for those who would like to explore the world of knocking on the registers of the Z8530 SCC.

Hint #1: Synchronize your code with the SCC logic.

Before writing to the SCC chip for the first time, you should make an attempt to ensure your code is synchronized with the SCC's logic. This needs to be done only once when you are initializing the SCC. This can be accomplished with a single read of SCC Read Register 0 (RR0). For example, if you're using serial port 2 (the modem port), the following code reads RR0 of SCC channel B:

```

longa off          must be using 8-bit accumulator
lda $C038         read RR0 of SCC Channel B

```

Hint #2: Watch out for interrupts from the other SCC channel.

Except for RR0, WR0, and the two SCC data registers, all SCC registers are accessed in a two-step process. First, the register number you want to select is written to WR0. After the register number is set, the next read from or write to the command register accesses the register selected in the first step. Because several of the SCC registers are shared between the two SCC channels and because code accessing them may not always be yours (i.e., AppleTalk), interrupts should be disabled during the two steps. The following code shows two quick subroutines to access the SCC's Read and Write registers while preventing interrupts between the register number set and the register read or write steps:

```

                                longa off          must be using 8-bit accumulator
                                longi off          and index registers
;
; Write to a SCC command register - channel A or B.
; Input:  A = value to store
;         X = SCC register number ($0-$F)
;         Y = $01 channel A
;         $00 channel B
;
WriteSCC                        php                save the current interrupt status
                                sei                disable interrupts
                                pha                save value to write
                                txa                get SCC register number from X
                                sta $C038,y       set the register number
                                pla                restore value to write
                                sta $C038,y       write the value
                                plp                restore the interrupt status
                                rts
;
; Read from a SCC command register - channel A or B.
; Input:  A = SCC register number ($0-$F)
;         Y = $01 channel A
;         $00 channel B
; Output: A = register value
;
ReadSCC                         php                save the current interrupt status
                                sei                disable interrupts
                                sta $C038,y       set the SCC register number
                                lda $C038,y       get the value from the SCC register
                                xba                look ahead 2 lines...
                                plp                restore the interrupt status
                                xba                set N and Z flags for exit
                                rts

```

Just to be complete, here's how RR0, WR0, the receive buffer, and the transmit buffer SCC registers are accessed on the Apple IIgs:

```

longa off          must be using 8-bit accumulator

```

```

                                longi off                                and index registers
;
; Read RR0 - channel A or B
; Input:  Y = $01 channel A
;          $00 channel B
; Output: A = RR0 register value
;
ReadRR0                          lda $C038,y          get the value from RR0
                                rts
;
; Write WR0 - channel A or B
; Input:  A = value to store at WR0
;          Y = $01 channel A
;          $00 channel B
;
WriteWR0                          sta $C038,y        write the value to WR0
                                rts
;
; Read from SCC receive buffer - channel A or B
; Input:  Y = $01 channel A
;          $00 channel B
; Output: A = value of data received
;
ReadData                          lda $C03A,y        get the value from SCC data register
                                rts
;
; Write to SCC transmit buffer - channel A or B
; Input:  A = value of data to transmit
;          Y = $01 channel A
;          $00 channel B
;
WriteData                          sta $C03A,y      write the value to SCC data register
                                rts

```

Hint #3: All SCC channels are not created equal.

In the IIGs, the SCC's receive and transmit clocks for both channels are driven by a single crystal oscillator circuit. This is accomplished by connecting a 3.6864 MHz crystal between the /RTxC and /SYNC pins of channel A. Channel B's /RTxC pin is connected to Channel A's /SYNC pin to drive channel B's clocks from channel A's oscillator circuit.

Because of this single circuit, Write Register 11 (WR11) bit 7 must be set to 1 for SCC channel A and must be set to 0 for SCC channel B.

Hint #4: RR3 is available only in SCC channel A.

When your interrupt handler is checking to see if the interrupt condition was caused by your SCC channel, remember to always look at RR3 in SCC channel A. RR3 in channel A contains the interrupt pending bits for both SCC channels. RR3 in channel B always returns all zeros, which doesn't tell you a lot about what's happening.

Don't be a Serial Killer

How to Install and Remove your SCC Interrupt Handler

If you're going to handle serial I/O and don't want your application to have to poll the SCC chip all the time to see if something has happened, you probably want to install an interrupt handling routine that is called every time a SCC chip condition you want to know about occurs. This section of the

Note shows how to install and remove your own SCC interrupt handler.

The steps for installing your SCC interrupt handler are:

1. Ensure the serial firmware's Input and Output buffering is disabled. The state of I/O buffering can be checked by looking at bit 14 of the ModeBitImage parameter returned by the GetModeBits extended interface call. I/O buffering can be disabled with the firmware's BD control command.
2. Disable the SCC Master Interrupt Enable (WR9, bit 3) briefly while performing the next six steps. The value you should write to WR9 is 00000010.
3. Get the address of the system interrupt flag byte, SerFlag. The ROM version determines the method of finding the address of SerFlag. In ROM version 01 and later, you can get the address with a call to the Miscellaneous Tools GetAddr using a reference number of \$000E. With ROM version 00 (the original IIgs ROM), the address of SerFlag is \$E10104. Refer to the Apple II Miscellaneous Technical Note #7, Apple II Family Identification for information on identifying Apple IIgs ROM versions.
4. Once you have the correct address of SerFlag, preserve the byte's current value, then turn on the bits in the byte which reflect the port from which you are handling interrupts. The bits for the different ports are as follows (note the relationship of the bits of RR3 to SerFlag):

```
Port 1:  ORA    %#00111000
Port 2:  ORA    %#00000111
```

5. Initialize the SCC modes. The Z8530 Serial Communications Controller Technical Manual shows the order the SCC registers must be programmed. However, you must stray from the manual slightly due to the hardware implementation of the SCC in the IIgs. A typical initialization sequence to set the SCC up for asynchronous serial communications through channel B (the modem port) would look similar to the following:

SCC Register	Value	Comment
RR0	-	ensure synchronization with SCC
WR4	01000100	x16 clock, 1 stop, no parity
WR3	11000000	8 bit receive data, auto enables off, receiver disabled
WR5	01100010	DTR is active, 8 bit transmit data, no break, transmit disabled, RTS is inactive
WR11	01010000	no Xtal on channel B, receive and transmit clock = baud rate generator output
WR12	01011110	low byte of baud rate generator time constant = \$5E - 1200 baud
WR13	00000000	high byte of baud rate generator time constant = \$00 - 1200 baud
WR14	00000000	no local loopback or auto echo, /DTR follows inverted DTR bit in WR5, use /RTxC for baud rate generator clock, disable baud rate generator
WR14	00000001	enable the baud rate generator
WR3	11000001	receiver enabled
WR5	01101010	transmit enabled
WR15	00000000	no interrupts on this channel for now...

6. Tell the SCC which external and status conditions can cause an interrupt by setting the appropriate bits in WR15. This step is

- not needed unless you are setting bit 0 of WR1 (External/Status Master Interrupt Enable) in the next step.
7. Enable the interrupts modes you want by setting the appropriate bits in WR1 (00010011 for all SCC interrupt conditions).
 8. Use `ALLOC_INTERRUPT` to add your interrupt handler to the operating system's interrupt vector table. The interrupt identification number returned by `ALLOC_INTERRUPT` is needed when you remove your interrupt handler.
 9. Reenable the SCC Master Interrupt flag (WR9, bit 3). The value you should write to WR9 is 00001010.

The interrupt handling routine must conform to the rules listed in the ProDOS 8 Technical Reference Manual and GS/OS Reference, Volume 2.

When you get ready to shut down your application, you need to remove your interrupt handler. The steps for removing the SCC interrupt handler you installed are as follows:

1. Disable the SCC Master Interrupt Enable (WR9, bit 3) briefly while performing the next six steps. The value you should write to WR9 is 00000010.
2. Disable all interrupts modes for your port by writing a \$00 to WR1.
3. Remove any character that might be left in the receive data register by reading it once.
4. Clear any pending transmit overrun and external and status interrupts by writing 11010000 to WR0.
5. Clear any pending transmit interrupt by writing 00101000 to WR0.
6. Use `DEALLOC_INTERRUPT` to remove your interrupt handler from the operating system's interrupt vector table.
7. Restore `SerFlag` to its original value.
8. Reenable the SCC Master Interrupt flag (WR9, bit 3). The value you should write to WR9 is 00001010.

Further Reference

-
- o Apple IIgs Toolbox Reference Manual, Volume 1
 - o Apple IIgs Firmware Reference Manual
 - o Apple IIgs Hardware Reference Manual, Second Edition
 - o GS/OS Reference, Volumes 1 and 2
 - o ProDOS 8 Technical Reference Manual
 - o Apple II Miscellaneous Technical Note #7, Apple II Family Identification
 - o GS/OS Technical Note #9, Interrupt Handling Anomalies
 - o Z8530 Serial Communications Controller Technical Manual (Zilog Corporation)
 - o Z85C30 Serial Communications Controller Technical Manual (Advanced Micro Devices, Inc.)

Apple IIGS

#19: Multichannel Output with the Apple IIGS Note Synthesizer

Revised by: Jim Mensch

November 1988

Written by: John Worthington & Jim Merritt

June 1987

This Technical Note discusses multichannel sound with the IIGS Note Synthesizer.

It is possible to play multichannel sound using the IIGS Note Synthesizer Tool Set. The Ensoniq Digital Oscillator Chip (DOC) supports 16 independent output channels. Since only the low three bits of the output channel number are available through the IIGS sound expansion connector, multichannel circuitry may only decode eight output channels (zero through seven). Output channel eight maps onto channel zero, channel nine onto channel one, etc., and this mapping continues through all 16 channels.

The setting of the high nibble of the DOCMode byte in a waveform of the waveList portion of the instrument definition determines the routing of output from a Note Synthesizer instrument to a particular channel (the actual DOCMode information is in the low nibble of the DOCMode byte). You may assign each separate element in a waveList to a different output channel to create multisampled instruments in which some samples play on the left speaker and others on the right.

Apple standards require stereo expansion cards to map all even output channels to the right and odd channels to the left. To be compatible with cards that decode more than two of the chip's output channels, software should use channel zero for right and channel one for left. This convention ensures that output is always positioned properly in the stereo space with channel zero information going to the right front and channel one information going to the left front.

Further Reference

- o Apple IIGS Toolbox Reference, Volume 2
- o Apple IIGS Toolbox Reference Update

Apple IIGS

#20: Catalog of APW Language Numbers

Revised by: Matt Deatherage

March 1990

Written by: Jim Merritt

August 1987

This Technical Note formerly listed APW Language Number assignments, which correspond to auxiliary type values of file type \$B0. Changes since November 1988: This information is now documented in Apple II File Type Notes, specifically Notes of file type \$B0.

The correspondence between APW Language Numbers and auxiliary type values for \$B0 files is no longer one-to-one. Although all APW Language Numbers are stored with their source files in the auxiliary type field, there now exist assignments of auxiliary type values for file type \$B0 which are not APW languages.

Therefore, the contents of this Note can now be found in the File Type Note for file type \$B0, where all such assignments of either kind are still called "APW Language Numbers."

Further Reference

- o File Type Note for file type \$B0, Apple IIGS source code files

Apple IIGS

#21: DMA Compatibility for Expansion RAM

Revised by: Glenn A. Baxter
Written by: Jim Merritt

November 1988
August 1987

This Technical Note discusses the Apple IIGS Extended Memory Slot specification.

The Apple IIGS Extended Memory Slot specification provides for DMA access to no more than four rows of RAM on a single board through the CROW0 and CROW1 signals. Expansion board designs that involve more than four rows of RAM are not compatible with DMA accesses. Each of the four rows can hold either 256K or 1 MB of data. The design of the Fast Processor Interface (FPI) imposes this limit. Each row can be organized in any of the following configurations to yield the respective board capacities assuming there are no more than four rows:

Chips	Configuration	Board Capacity
8	256K x 1 DRAM	1 MB
8	1 MB x 1 DRAM	4 MB
2	256K x 4 DRAM	1 MB
2	1 MB x 4 DRAM	4 MB

The CROW0 and CROW1 signals properly decode the row addresses for both normal and DMA cycles. The Extended Memory Slot interface does not support the latching of bank address information off the data bus during a DMA cycle, and a card which attempts to latch the bank address will likely get the last CPU cycle's bank address. Getting the last address is not a problem if it accidentally happens to be the bank to which you wish to talk, but this is rarely the case. The card gets the last CPU cycle's bank address because DMA essentially shuts off the CPU, so it cannot emit the bank address. The FPI, which contains the DMA bank address register (\$C037), does not emit the DMA bank address either, thus preventing bus contention with the processor as it is being removed from that bus. The DMA bank address register inside the FPI affects the addressing and control information that the Extended Memory Slot sees; it does not affect the data bus. Therefore, during DMA, the bank address time is filled with what is essentially random bank address information. Using this random information could result in damaging the contents of the memory (destroying little things like the operating system).

Suppose a card were designed to latch the bank address directly from the data bus with the rising edge of the PH2 clock signal. It could use the bank address to derive the proper RAM row address and never bother with CROW0 and CROW1 at all. Directly latching the bank address would permit the card to accommodate any desired RAM arrangement in 64K increments, including an odd number of rows. Although the technique is valid during CPU cycles, it does not work during DMA cycles since the FPI never emits the DMA bank address onto the data bus. During DMA cycles, any card that tries to latch the bank address directly, instead latches the bank address that was put on the data bus during the last CPU cycle, which is probably the wrong value.

Currently, there does not seem to be a solution for the DMA situation. There is the possibility of "limited DMA compatibility." An example of a limited-compatibility card would be one with six banks of memory. Its lower four

banks are DMA compatible since they use the CROW0 and CROW1 lines, but the upper two banks do not work properly with DMA. This limited approach should be safe, but it is not guaranteed since DMA cards are sometimes aware of the total system memory and may expect, quite reasonably, to have access to all of the memory when in fact it does not. There are currently no "memory intelligent" DMA cards, but that could change at any point. The best we can suggest at this time is for hardware developers to build only four-row cards allowing up to 4 MB of memory, which is sufficient for most current applications.

Further Reference

- o Apple IIGS Hardware Reference

Apple IIgs
#22: Proper Use of Dynamic Segments

Rewritten by: Eric Soldan & Andy Stadler
Written by: Guillermo Ortiz

September 1990
October 1987

This Technical Note discusses strategies that applications can use to deal with dynamic segments.
Changes since November 1988: Rewrote from scratch to address current problems.

When reading the documentation on dynamic segments, it initially appears that they are even better than sliced bread. While they are incredibly useful, there are two issues that make dealing with them somewhat tricky. The first involves loading a dynamic segment; the second involves unloading a dynamic segment. Everything else works fine.

Loading Dynamic Segments

Loading dynamic segments is supposed to happen automatically. You are supposed to be able to call the code in the dynamic segment, and the system automatically loads it. As long as there is enough RAM to load the segment, this is exactly what happens.

The problem arises when there isn't enough memory. Immediately you have a number of questions, such as "How do I know if it didn't load?" and "How is the not-enough-memory error returned?" Unfortunately, neither of these questions is applicable. Instead, you get a Fatal System Error, which is not the most useful thing that could happen.

However, there are some reasons for this error. For example, in the Pascal or Toolbox stack frame system, the called function is responsible for removing the parameters pushed onto the stack. If the dynamic segment did not load, these parameters cannot be pulled from the stack, and if they are not pulled from the stack, the operating system cannot return to the caller.

Due to this problem, the best thing to do is to try to load the dynamic segment with LoadSegName. If it loads, then there is (obviously) enough RAM for it. If it does not load, then there was not enough RAM; it's that simple. So, to call a function named dynFN in a dynamic segment called dynSeg, you would do the following:

```
LoadSegName("\pDynSeg");
if (!_toolErr) {
    dynFN(some, number, of, parameters);
    UnLoadSeg(dynFN);
}
else ErrorAlert("\pOut of RAM.");
```

Unloading Dynamic Segments

UnLoadSeg used to have a problem, so the above technique would not have worked. As of System Software 5.0.3, this problem has been fixed. In the example, the code UnLoadSeg(dynFN) does not pass the address of the dynFN that

was loaded into RAM. Instead, that address represents the entry in the dynamic segment jump table for that particular function. The jump table is always in RAM. So, you are not actually passing an address of the segment to be unloaded, but an address in the jump table.

The loader is responsible for figuring out that the address is actually an address in the jump table, and it is supposed to unload the segment to which the jump table entry refers. The loader did not handle this case properly until 5.0.3. So, for system disks prior to System Disk 5.0.3, you can preserve the segment number returned by the LoadSegName call to issue an UnLoadSegNum call to dispose of the dynamic segment. Due to UnLoadSeg not doing the job prior to 5.0.3, you could use UnLoadSegNum. This also has problems. ExpressLoad changes the segment numbers, so it is difficult to maintain the segment numbers if you change the link script. For these reasons, the below technique should be used for system disks prior to 5.0.3:

```
void sample()
{
    struct LoadSegNameOut dynSegInfo;

    dynSegInfo = LoadSegName("\pDynSeg");
    if (!_toolErr) {
        dynFN(some, number, of, parameters);
        UnLoadSegNum(dynSegInfo.segNum);
    }
    else ErrorAlert("\pOut of RAM.");
}
```

Dynamic Segment Interdependencies: Just Say No

Dynamic Segments calling each other almost always lead to unloading conflicts, and more importantly, they defeat the purpose (if they both have to be in simultaneously then they might as well be static). Figure 1 is a sample program layout you may want to consider when designing your application dynamic segment usage:

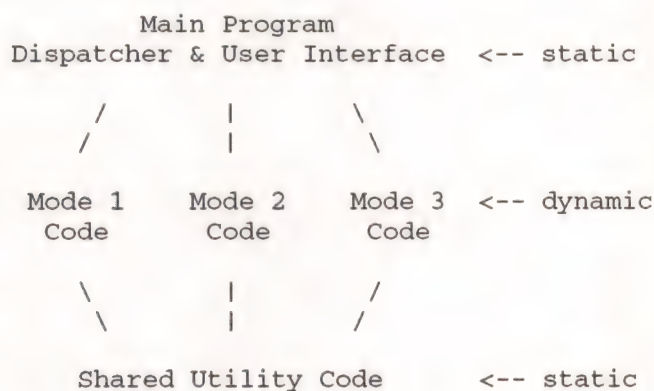


Figure 1-Sample Program Layout

Also, if one of the dynamic segments described is much more than, say, 32K or 40K, you may wish to load a pair (or more) of dynamic segments. These dynamic segment pairs would always be loaded and unloaded simultaneously. Why? Because loading two 25K segments is more likely to succeed than loading one 50K segment.

A Final Warning:

Data in a dynamic segment is a tricky issue. When you call a dynamic segment, you are not sure if it got loaded, or if it was already in RAM, and therefore you cannot be sure of the values in your global data. For example, say that you have a global variable that represents the number of times that you call the dynamic segment. Every time you call the segment, you would increment this variable. This technique works great until the dynamic segment gets purged. Once it is purged, the next time you call it, the variable area would be loaded from disk again, with its original initial value. The count is no longer valid. To fix this, you can place the global count variable in the static globals space for the main code. Then the variable would not get purged, and your count would be valid. Of course, if you have global data that does not ever change, then it is okay for the data to be in the global segment.

Further Reference

- o GS/OS Reference
- o Apple IIgs Programmer's Workshop Assembler Reference

Apple IIGS
#23: Toolbox Use of DOC RAM

Revised by: Matthew Denman & Matt Deatherage
Written by: Jim Merritt

November 1988
October 1987

This Technical Note explains why you must be careful about which values you store in the first page of the Ensoniq Digital Oscillator Chip (DOC) RAM when using Note Synthesizer and MIDI Tool Sets on the Apple IIGS.

The Apple IIGS Note Synthesizer uses an oscillator as a free-running timer to clock the update of waveform envelopes when the DOC sounds notes. To act as a timer, the oscillator "plays" the contents of bytes \$00 - \$FF in DOC RAM at zero volume. Once it scans through the entire "waveform buffer," the oscillator generates an interrupt, which the appropriate Note Synthesizer routines service.

When using the Note Synthesizer or the Note Sequencer without the MIDI Tool Set, there is no need to avoid using DOC RAM locations \$00 - \$FF for general waveform storage. More than one oscillator can play from the same waveform buffer at the same time, so the function of the timer oscillator does not affect normal use of the DOC for sound generation purposes in any way. However, you should not fill the first page of DOC RAM with waveforms that are delimited by zero bytes (as is sometimes appropriate in special situations, discussion of which is beyond the scope of this Note). The presence of zero bytes in the first page of DOC RAM can cause serious system performance degradation and can even cause the system to hang. In particular, it is always inappropriate to store arbitrary, non-waveform data in the first page of DOC RAM since such data often includes zero bytes (which would be corrupted were you to remove or modify them).

The Apple IIGS MIDI Tool Set also uses bytes \$00 - \$FF of DOC RAM for timing purposes, but it uses a different oscillator than the Note Synthesizer. If you want MIDI time stamping, you may not use the first page (bytes \$00 - \$FF) of DOC RAM for your own purposes since the MIDI Tool Set uses the contents of those bytes for time-stamping purposes.

You may use the MIDI, Note Synthesizer, and Note Sequencer Tool Sets together, but you must not use bytes \$00 - \$FF of DOC RAM for any purpose if using MIDI time stamping, nor store zero bytes in this area when using the Note Synthesizer. You might consider it appropriate to avoid using the first page of DOC RAM, if possible, to facilitate adding MIDI support to your application at a later date.

Apple IIgs

#24: Apple IIgs Toolbox Reference Updates

Revised by: Dave Lyons

May 1992

Written by: Rilla Reynolds, Matt Deatherage, Dave Lyons,
C. K. Haun & Eric Soldan

October 1987

This Technical Note documents changes to the Apple IIgs Toolbox Reference manuals. Please contact Apple II Developer Technical Support at the address listed in Apple II Technical Note #0 if you have additional corrections or suggestions for any of the Apple IIgs Toolbox documentation.

CHANGES SINCE DECEMBER 1991: Added corrections to Dialog Manager, Menu Manager, Tool Locator, Window Manager, and Appendix E.

The current Apple IIgs Toolbox reference material is Apple IIgs Toolbox Reference, volumes 1 to 3 as well as this Technical Note. (The Apple IIgs Toolbox Reference Update beta draft from APDA is obsolete and should not be used.)

CORRECTIONS TO VOLUME 1

DESK MANAGER--FIXAPPLEMENU CAN DIE WITH ERROR \$0512

Fatal system error \$0512 comes from FixAppleMenu (in the Desk Manager). It means that one of your installed New Desk Accessories does not have a well-formed menu title string. In particular, the required backslash (\) character was not found (make sure bit seven is off).

DIALOG MANAGER--EDITLINE ITEM VALUE

On page 6-12, the description of an editLine item value should read "Maximum length of the item text (0 to 255 characters)."

THE LIST MANAGER WANTS THE PORT SET PROPERLY

The List Manager expects the current grafPort to be set properly before you make most List Manager calls; drawing can occur in funny places if the grafPort is not set properly before calls that draw (like SelectMember2). Most List Manager calls, and many other toolbox calls, require that the current grafPort be explicitly set. Before you call List Manager routines that draw, set the current port to your window with a SetPort call. Remember the note in Volume 2 under the NewWindow call--"Important: NewWindow does not set the current port, but many routines require that a current port exist. Use the QuickDraw II routine SetPort to set the current port." Using SetPort can prevent toolbox confusion and reduce your debugging time.

DELETITEM OPERATES ON THE CURRENT MENU BAR

Page 13-37 says DeleteMItem removes the specified item from the current menu. It means the item is removed from the current menu bar.

ERROR \$0F02 FROM GETMITEM

GetMItem returns error \$0F02 if the specified menu item is not found.

On page 13-45, the return value from GetMenuFlag should read "Word--menuFlag value for the specified menu."

On page 13-56, in the description of the hiliteFlag parameter to HiliteMenu, no particular value of "TRUE" is specified. \$0001 is a good value (\$8000 does not work; bit 15 is special).

On page 13-72, SetMenuFlag doesn't bother to actually explain what it does. If bit 15 of newValue is zero, each set bit set forces the corresponding bit in the menu's flag value to be set. If bit 15 of newValue is one, each clear bit forces the corresponding bit in the menu's flag value to be clear. Knowing this, you can set or clear more than one bit at a time, if you want.

SETVECTOR REFERENCE NUMBERS

On page 14-62, vector reference number \$002C is listed as "Message pointer vector." \$002C is actually the stack-based GS/OS call vector. (The real message pointer vector is not accessible through GetVector and SetVector.)

GETTING A CLEAN MOUSE MODE FROM READMOUSE

On ROM 3 computers, the mouse mode byte returned from ReadMouse sometimes has extra bits set in the high nibble. Before feeding a ReadMouse value to SetMouse, mask off all but the low nibble (AND #\$000F).

READASCIITIME RESULT BUFFER

The description of ReadAsciiTime (in the Miscellaneous Tools) on page 14-16 should say the most significant bit (not byte) of each character is set to one.

SYSTEMEVENT IS ALL BACKWARDS

Although applications still should not call SystemEvent, we should note for completeness that the input parameters listed in Volume 1 are exactly backwards in the stack diagram.

CORRECTIONS TO VOLUME 2

QUICKDRAW AUXILIARY ERROR CODES

Following are some error codes from QuickDraw Auxiliary that are not listed in volume 2.

\$1210: picEmpty
\$1211: picAlreadyOpen
\$1212: pictureError

\$1221: badRect
\$1222: badMode

FRAMERGN DOES NOT CONTRIBUTE TO AN OPEN REGION

The description of the FrameRgn routine on page 16-105 in the Apple IIgs Toolbox Reference, Volume 2 states that FrameRgn will contribute to a region definition if a region is open when FrameRgn is called. This is incorrect; FrameRgn does not contribute to the region being defined. To add a region to another region, use XorRgn or UnionRgn.

TOOL LOCATOR, TLMOUNTVOLUME

On page 24-21, the description of TLMountVolume does not bother to mention that QuickDraw II and Event Manager must be active. If they are not, you should use TLTextMountVolume instead.

TOOL LOCATOR, SETTSPTR

When using SetTSPtr to patch a system tool set, the Tool Locator and Desk Manager are special. See Apple IIgs Technical Note #101, Patching the Toolbox.

WINDOW MANAGER, "DRAW INFORMATION BAR ROUTINE"

On page 25-23, the code to clean up the stack is incorrect. On the sta <14, the comment "Works because stack and direct page are equal" is no longer true--they were equal until the PLY two lines earlier. One way to correct the code is to replace sta <14 with sta 14,s and sta <12 with sta 12,s.

WINDOW MANAGER, INVALRECT

The description of InvalRect on page 25-80 claims that InvalRect modifies the input rectangle; the rectangle is actually not modified.

WINDOW MANAGER, PINRECT

On page 25-89, in the description of PinRect, the two greater-than comparisons should be greater-than-or-equal.

WINDOW MANAGER, SETZOOMRECT

The description of SetZoomRect on page 25-112 refers to fZoomed as bit 2 in the window frame. fZoomed is actually bit 1, with value \$0002.

WINDOW RECORD OFFSETS

On page 25-142, note that the offsets given into the window record refer to the record as the Window Manager treats it internally, with a wNext field at the beginning. When dealing with a window pointer as seen by an application, you need to subtract four from the offsets shown. For example, wPort is \$00 (not \$04), and wControls is \$C6 (not \$CA).

APPENDIX A, "WRITING YOUR OWN TOOL SETS"

At the bottom of page A-8, "lda #\$90" should read "lda #\$8100" for version 1.0 prototype.

On page A-10, the figure should show two RTL addresses (6 bytes) on the stack.

CORRECTIONS TO VOLUME 3

CONTROL MANAGER: MENU EVENTS

On page 28-15, note that a Menu Event is identified by the value wInSpecial (\$0019) in the what field of the task record. The menu item ID is in the low word of the wmTaskData field.

CONTROL MANAGER: DIMMED CUSTOM CONTROLS

In the Draw routine for both extended and non-extended controls, the high word of ctlParam (which was previously undocumented) contains a flag which the definition procedure can use to draw a normal or dimmed control. The value is \$0000 normally, but it is \$FFFF when the control is inactive (hilite value

equals \$00FF), or when the control's state is tied to the window's state and the window is inactive.

CONTROL MANAGER: SIZE BOX CONTROLS

The part code for an extended Size Box control is normally 10. If the `fCallWindowMgr` bit is set in `ctlFlag`, the part code is \$80; and if the size box is managed by a Text Edit control, the part code is \$84.

When a Size Box control's `fCallWindowMgr` bit is set, the control needs to pass a minimum window size to `GrowWindow`. It gets this value from its `ctlData` field, which you can get with `GetCtlTitle` and set with `SetCtlTitle` (the low word is the minimum height, and the high word is the minimum width). A height of zero defaults to 50, and a width of zero defaults to 130.

DESK MANAGER: ERRORS FROM ADDTORUNQ AND REMOVEFROMRUNQ

The Desk Manager chapter, page 29-6, states no errors are possible for `AddToRunQ`, but any errors from the Miscellaneous Tools routine `AddToQueue` are returned unchanged.

Page 29-8 states no errors are possible from `RemoveFromRunQ`, but any errors from `DeleteFromQueue` are returned unchanged.

EVENT MANAGER: WHAT SETAUTOKEYLIMIT REALLY DOES

Page 31-6 says that `PostEvent` will add up to the new auto-key limit number of auto-key events before reverting to the rule that auto-key events are only to be posted if the event queue is empty. This is not quite right. Actually, the parameter to `SetAutoKeyLimit` is used in a size comparison on the event queue--if there are `newLimit` or more events in the queue, auto-key events will not be posted. Volume 3 incorrectly states that up to `newLimit` auto-key events will be posted; this is only true if you assume the event queue is empty before the first auto-key event comes in.

LIST MANAGER

On page 35-9, the description of `ResetMember2` does not point out an important difference between `ResetMember2` and `NextMember2`. `ResetMember2` deselects the member found, but `NextMember2` does not change the member's status.

On page 35-3, bit 5 of the `memFlag` field is defined--it makes an item inactive. To make use of this bit, you must also set bit 6 of the List control's `ctlFlag` field; if you don't set this bit, the user will still be able to select members using the mouse.

MEMORY MANAGER

If the Memory Manager detects a corrupted entry in the Out Of Memory Queue, fatal system error \$0209 occurs.

MENU MANAGER

On page 28-65, the description of the `initialValue` field is misleading. Cross out the text "that is, its relative position within the array of items for the menu." `initialValue` is simply a menu item ID, not an offset into an array.

Page 37-7 states "Because caching does not work with menus in windows, the `InsertMenu` call automatically disabled caching for such menus." Actually, `InsertMenu` doesn't do that. You should not set the `allowCache` bit for a menu in a window.

MISCELLANEOUS TOOLS: INTERRUPT STATE RECORD NOT ALWAYS COMPLETE

The interrupt state record returned from GetInterruptState (and passed to SetInterruptState) is not always completely filled in. The Interrupt Manager, in the interest of serving AppleTalk and serial interrupts as rapidly as possible, does not take the time to save all the items in the record until those timing-critical interrupt handlers have been called. Some items are not saved at all unless the interrupt is determined to be a BRK instruction. Table 1 shows all items in the current interrupt state record and when they become valid:

Record variable	When valid
irq_A	always
irq_X	always
irq_Y	always
irq_S	after serial
irq_D	always
irq_P	only on break
irq_DB	after serial
irq_e	after serial
irq_K	only on break
irq_PC	only on break
irq_state	after serial
irq_shadow	always
irq_mslot	after serial

Table 1--Validity of Interrupt Record

STANDARD FILE

On page 48-39, the description of origNameRef reads "On output, this string contains the string confirmed by the user, which may not be the same length as the default value." This sentence is confused; ignore it. The string is not changed at all; Standard File doesn't even know how long the buffer is.

TOOL LOCATOR: NOTES ON STARTUPTOOLS

StartUpTools in System Software 5.0.4 and earlier is intended to be used from applications only, not from NDAs.

The order of the toolArray entries in the StartStop record is not important. StartUpTools and ShutDownTools always start up and shut down tools in a correct order.

StartUpTools in System Software 5.0.4 and earlier fails to open your application's resource fork if the application's filename contains a slash (/) or if the application directory path is longer than 64 characters.

For maximum compatibility, pass your application's master user ID with any auxID to StartUpTools instead of allocating a new user ID.

WINDOW MANAGER:NEWWINDOW2 PARAMETERS OVERRIDE TEMPLATE EVEN WHEN YOU PASS NIL

The description of the NewWindow2 call on page 52-32 is in error. The description of the titlePtr, refCon, contentDrawPtr, and defProcPtr says, "To prevent NewWindow2 from replacing the template values, supply NIL pointers..." This is only true for the titlePtr parameter--if you pass NIL for any of the other parameters then the value of that parameter in your window record is also NIL, no matter what the template value was. In other words, if you have the value \$99 stored in your template refCon field, and you pass NIL for the

refCon value in a NewWindow2 call, the value of the refCon in the returned grafPortPtr is NIL.

APPENDIX E: RTEXTFORLETTEXTBOX2 RESOURCES

Page E-68 of Volume 3 shows a length field at the beginning of an rTextForLETextBox2 resource. This field is not actually present. The length is simply the size of the resource--it is not stored redundantly.

APPENDIX E: RTWORECTS RESOURCES

When the two rectangles are for 320- and 640-mode, by convention the rectangle for 320 mode comes first.

Further Reference:

- o Apple IIgs Toolbox Reference, Volumes 1-3
- o Apple IIgs Technical Note #101, Patching the Toolbox

Apple IIgs
#25: Apple IIgs Firmware Reference Updates

Revised by: Dave Lyons
Written by: Rilla Reynolds, Dave Lyons
& Jim Luther

May 1992
October 1987 to September 1990

This Technical Note includes updates to the May 1987 edition of the Apple IIgs Firmware Reference, published by Addison-Wesley (Part Number 030-3121-A). The new Monitor commands require an Apple IIgs revised ROM (Part Number 342-0077-B), which is available without charge from an authorized Apple dealer. Please contact Apple II Developer Technical Support at the address listed in Apple II Technical Note #0 if you have additional corrections or suggestions for this manual.

CHANGES SINCE SEPTEMBER 1990: Added a reference to Apple IIgs Technical Note #102 for TOBRAMSETUP.

CONTENTS

Page vii, Chapter 7 SmartPort Firmware: Change "Generic SmartPort calls 121" to "Standard and Extended SmartPort calls 121."

CHAPTER 2: NOTES FOR PROGRAMMERS

Page 11, Environment for the Firmware Routines: Refer to Apple IIgs Technical Note #88, The Page One Stack in a 16-Bit World for more information on manipulating the stack pointer.

CHAPTER 3: SYSTEM MONITOR FIRMWARE

Page 24, Table 3-1 (continued), Monitor commands grouped by type: Add a miscellaneous-type and a debugging-type Monitor command to the table, as follows:

Command type	Command format
...	
Quit Monitor	Q
Install Visit Monitor and MemoryPeeker desk accessories	#
...	
Enter mini-assembler	!
Set flags (e, m, x) for full-native mode	Control-N

Page 43, Back to BASIC: The last paragraph should read: "If you are using DOS 3.3 or ProDOS(R), use the Monitor Q (Quit) command to return to the language you were using with your program and variables intact."

Page 48, Table 3-6, Commands for program execution and debugging: Add a Monitor command to the table:

Command type	Command format
...	
Enter mini-assembler	!
Set flags (e, m, x) for full-native mode	Control-N

Page 66, after final paragraph: Add a new Monitor instruction heading and description:

NATIVE MODE SET CONTROL-N (NATIVE MODE)

Control-N sets the m, x, e flags to 0 for full-native mode. All other registers are unchanged.

Page 67, after final paragraph: Add a new Monitor instruction heading and description:

TURN ON ROM DESK ACCESSORIES, #

Enables the currently available ROM desk accessories, Visit Monitor and Memory Peeker. These desk accessories remain active in the desk accessory menu until power is shut off. Control-Open Apple-Reset has no affect on these items. To exit the Visit Monitor desk accessory, press Control-Y then press Return. To exit the Memory Peeker desk accessory, press Q.

CHAPTER 4: VIDEO FIRMWARE

Page 77, Table 4-4, Control characters with 80-column firmware on: Change the actions taken by Control-E and Control-F to read (they were reversed):

Control character	Action taken by C3COUT1
Control-E	Turns cursor on
Control-F	Turns cursor off

CHAPTER 5: SERIAL-PORT FIRMWARE

Page 82, Compatibility: The second half of the third sentence in the first paragraph should read: "...the Apple IIgs hardware is different from that used on the SSC."

Page 91, Input buffering, BE and BD: This heading should be "Input/Output buffering, BE and BD."

Page 94, Table 5-6: The Extended Interface footnote which states, "If the function call returns with the carry bit set..." is incorrect. For Apple IIgs ROM 01, the Extended Serial Interface does not return the error condition in the carry bit. Programs using the Extended Serial Interface should check for a non-zero result value in the result code rather than the carry bit to determine if an error has occurred. For additional error handling information using the Extended Interface, see Apple IIgs

Technical Note #50, Extended Serial Interface Error Handling.

Page 95, Error handling: The second sentence should read: "If the character has a framing or parity error (assuming that the parity option is not set to None), the character is deleted from the input stream and the appropriate mode bit is set."

Page 96, Note: The Note should read: "The InQStatus elapsed-time counter functions correctly only if a heartbeat interrupt task has been started. A heartbeat interrupt task is a set of functions called by interrupt code that run automatically at one-thirtieth of a second intervals."

Page 96, Interrupt notification: The fourth sentence in the first paragraph should be: "The system interrupt handler will transfer control to the user's interrupt vector at \$03FE in bank \$00."

Page 97, Interrupt notification: The last three paragraphs should be replaced with this paragraph: "The interrupt completion routine executes as part of the firmware interrupt handler and must be run in that environment. The interrupt completion routine must preserve the DBR, speed, 8-bit native mode, D register, stack pointer (or just use the current stack), and MSL0T for proper operation. A/X/Y need not be preserved."

Page 100, SetModeBits: The first sentence in the paragraph following the CMDLIST should read: "Use this call to alter any of the mode bits whose function is described below."

Page 105, GetIntInfo: The command list should read:

CMDLIST	DFB	\$03		;Parameter count
	DFB	\$0C		;Command code
	DW	\$00		;result code (output)
	DW	\$00		;interrupt setting (output)
	DL	Completion address		;(output)

The following should be added after the command list: "Note: The Parameter count of \$03 is correct even though there are four parameters."

The following should be added after the last paragraph: "Note: Before making this call from an interrupt completion routine, you must set the operating environment to look and act exactly like a 6502 in all respects. During interrupt completion routines, you must preserve the DBR, speed, 8-bit native mode, D register, stack pointer (or just use the current stack), and MSL0T for proper operation. A/X/Y need not be preserved. See "Environments for the Firmware Routines" in chapter 2, Notes for Programmers for details about setting and restoring the operating environment."

Page 106, SetIntInfo: The command list should read:

```

CMDLIST   DFB   $03           ;Parameter count
          DFB   $0D           ;Command code
          DW    $00           ;result code (output)
          DW    Interrupt setting ;(input)
          DL    Completion address ;(input)

```

The following should be added after the command list, "Note: The Parameter count of \$03 is correct even though there are four parameters."

CHAPTER 7: SMARTPORT FIRMWARE

Page 120, Issuing a call to SmartPort: The standard and extended SmartPort call examples should be:

This is an example of a standard SmartPort call:

```

SP_CALL   JSR   DISPATCH       ;Call SmartPort command dispatcher
          DC    i1'CMDNUM'      ;This specifies the command type
          DC    i2'CMDLIST'     ;Word ptr to param list in bnk $00
          BCS   ERROR           ;Carry is set on an error

```

This is an example of an extended SmartPort call:

```

SP_EXT_CALL JSR   DISPATCH       ;Call SmartPort command dispatcher
          DC    i1'CMDNUM+$40'   ;This specifies the ext cmd type
          DC    i4'CMDLIST'     ;Pointer to the parameter list
          BCS   ERROR           ;Carry is set on an error

```

Page 121, Generic SmartPort calls: Change occurrences of "Generic SmartPort Calls" to "Standard and Extended SmartPort Calls" in the header and the first sentence. Refer to SmartPort Technical Note #2, SmartPort Calls Updated, for updated information on the SmartPort STATUS call.

Page 122, Statcode = \$00: Change the function of bit 0 of the first device status byte to: "1 = Device currently open (character devices only) or disk switched (block device only)."

Page 124: SmartPort device types should be same as those documented in SmartPort Technical Note #4, SmartPort Device Types.

Page 125, SmartPort driver status: See SmartPort Technical Note #2, SmartPort Calls Updated, for the correct format of the status list for unit 0, status code 0.

Vendors must request a Vendor ID Assignment from Developer Technical Support before using a specific value in bytes two and three.

Page 125, Possible errors: Add the following:

\$1F No interrupt. Interrupts not supported.
 \$2B No write. Disk write-protected.
 \$2F Offline. Disk off-line or no disk in drive.

Page 126, ReadBlock: Add a sentence at the end of the first paragraph which reads, "On return, the X and Y registers indicate the number of bytes transferred."

Page 131, Open: The following changes apply for the CMDNUM:

	Standard call	Extended call
CMDNUM	\$06	\$46

Page 132, Read: Add a sentence at the end of the first paragraph which reads, "On return, the X and Y registers indicate the number of bytes transferred."

Page 140, Figure 7-8, Disk-sector format: Change to the following:

```

-----
|13      |F|D|A|9|T|S|S|F|A|D|A|F|1      |F|D|A|A|S|699      |4|D|A|F|
|5-Nibble|F|5|A|6|r|e|i|i|o|d|E|A|F|5-Nibble|F|5|A|D|e|GCR      |E|A|F|
|SelfSync| | | |a|c|d|r|r| | | |SelfSync| | | |c|Nibbles |C| | | |
|Fields  | | | |c|t|e|m|s| | | |Fields  | | | |t|Fields  |h| | | |
|         | | | |k|o| |a|L| | | |         | | | |o|         |e| | | |
|         | | | |r| |t|R| | | |         | | | |r|         |c| | | |
|         | | | | | | |C| | | |         | | | | | |         |k| | | |
|         | | | | | | | | | |         | | | | | |         |s| | | |
|         | | | | | | | | | |         | | | | | |         |u| | | |
|         | | | | | | | | | |         | | | | | |         |m| | | |
-----

```

A SelfSync Field is four 20 microsecond selfsync nibbles written as a sequence of five 16 microsecond nibbles.

Page 140, ResetHook: The Control code and Control list should be:

Control Code	Control list
\$06	Count low byte \$04 Count high byte \$00 Hook reference number \$xx, \$00, \$00, \$00

Page 141, SetInterleave: The Control code and Control list should be:

Control Code	Control list
\$0A	Count low byte \$01 Count high byte \$00 Interleave \$01 to \$0C

Page 143, UniDiskStat: The Status code and Status list should be:

Status Code	Status list

\$05	Byte	\$04
	Soft error	\$00
	Retries	\$xx
	A register after execute	\$xx
	Y register after execute	\$xx
	P register after execute	\$xx
	Byte	\$xx

Page 152, Passing parameters to a ROM disk: Add a sentence to the end of the second paragraph which reads: "These locations will not be preserved between SmartPort calls."

Page 156, Table 7-6, SmartPort error codes: Add the following error code:

Acc value	Error type	Description
\$69	IOTERM	I/O terminated due to new line

Page 166, Table 7-8, Standard command packet contents":

Byte 3 descriptions should read "Byte 2 of param list."
 Byte 4 descriptions should read "Byte 3 of param list."
 Byte 5 descriptions should read "Byte 4 of param list."
 Byte 6 descriptions should read "Byte 5 of param list."
 Byte 7 descriptions should read "Byte 6 of param list."
 Byte 8 descriptions should read "Byte 7 of param list."
 Byte 9 descriptions should read "Byte 8 of param list."

CHAPTER 8: INTERRUPT-HANDLER FIRMWARE

Page 184, Serial-port interrupt notification: The last three paragraphs should be replaced with this paragraph: "The interrupt completion routine executes as part of the firmware interrupt handler and must be run in that environment. The interrupt completion routine must preserve the DBR, speed, 8-bit native mode, D register, stack pointer (or just use the current stack), and MSL0T for proper operation. A/X/Y need not be preserved."

CHAPTER 9: APPLE DESKTOP BUS MICROCONTROLLER

Page 191, Sync, \$07: The first sentence should read: "This command performs the three preceding commands in sequence."

Page 194, Receive Bytes, \$48: The fourth sentence should read: "The second byte value is a combination of the device address in the high nibble and the ADB command in the low nibble (see the Apple IIgs Hardware Reference)."

CHAPTER 10: MOUSE FIRMWARE

Page 201: Mouse button positions should be changed as follows:

- o X data byte
 If bit 7 = 0, then mouse button 1 is down.

If bit 7 = 1, then mouse button 1 is up.

- o Y data byte
 - If bit 7 = 0, then mouse button 0 is down.
 - If bit 7 = 1, then mouse button 0 is up.

Page 205, Figure 10-1, Position and status information:

Bit 7 description should be: "Currently, button 0 is up/down (0/1)."

Bit 6 description should be: "Previously, button 0 was up/down (0/1)."

APPENDIX B: FIRMWARE ID BYTES

Page 223, Table B-2, Register bit information: Change the table to show that Bits 7-0 of the Y register hold the ROM version number, and the X register is reserved. In addition, the table description should read: "The Y register contains the machine ID and the ROM version number. The X register is reserved."

Page 249, COUT1: In the third sentence, change the value of line feed from \$8C to \$8A.

Page 277, RDALTZP: Change the comment to read: "Bit 7 = 1 if alt zp enabled."

APPENDIX D: VECTORS

Page 272: At the end of the introductory paragraph, add "The vectors TOWRITEBRAM through TOPRINTMSG8 must be called in eight-bit native mode."

See Apple IIgs Technical Note #102, Various Vectors, for more information about the TOBRAMSETUP vector.

Further Reference:

- o Apple IIgs Firmware Reference
- o Apple IIgs Firmware Reference 1MB Apple IIgs Update
- o Apple IIgs Technical Note #50, Extended Serial Interface Handling
- o Apple IIgs Technical Note #102, Various Vectors
- o SmartPort Technical Note #2, SmartPort Calls Updated

Apple IIGS
#26: ROM Revision Summary

Revised by: Matt Deatherage
Written by: Rilla Reynolds

September 1989
October 1987

This Technical Note summarizes revisions to the Apple IIGS ROM.
Changes since November 1988: Revised to cover ROM 3.

Apple currently supports two configurations of the Apple IIGS ROM, ROM 1 and ROM 3. In August 1989, Apple IIGS computers began shipping with a 256K ROM, referred to as version 3 or ROM 3 (ROM 2 was skipped since there was already enough confusion about the first version, ROM 0, and the second version, ROM 1). System Software continues to support ROM 1, but it no longer supports ROM 0. Authorized Apple dealers can upgrade older systems (i.e., machines with serial numbers lower than E704...) to ROM 1 upon request.

ROM 1 requires System Software 2.0 or later, while ROM 3 requires System Software 5.0 or later. Although applications may work using older system software releases, they may not function properly due to the coordination of system software and ROM revisions.

Changes from ROM 0 to ROM 1

ADB

- o Absolute ADB devices are now supported correctly.
- o ADB fatal system error code is now \$0911 instead of \$0400.
- o ADBReset routine now delays about 160 microseconds before reading the buttons.
- o ADBStatus TRUE is now \$FFFF instead of \$0001.
- o All ADB error codes now include the tool number.
- o SRQrmv no longer crashes when you make the call with a command pending.

AppleDisk 3.5

- o AppleDisk 3.5 Macintosh block reads and writes now work as documented.
- o Extended status call now returns bit 0 = 1 if AppleDisk 3.5 media has been switched since the last READ, WRITE, or FORMAT.
- o New AppleDisk 3.5 status calls have been implemented to get internal variable and work buffer starting addresses.

AppleTalk

- o Link Access Protocol (LAP) inter-packet gap now handles added SCC delay.
- o Name Binding Protocol (NBP) now considers uppercase and lowercase characters identical.
- o A nonexistent protocol no longer hangs the dispatcher.

Desk Manager

- o SaveScreen and RestoreScreen now work.

Event Manager

- o Now auto-key events are not posted in the queue unless the queue is empty.
- o EMStartUp and EMShutDown code has been optimized.
- o Event Manager now returns an error instead of crashing when there is an attempt to post an invalid event.

Integer Math

New Changes:

- o Optimized the multiply routine.

RAM patches moved to ROM:

- o Changes to FixMul, FixRatio, and SDivide.
- o SDivide recovers from a divide by zero operation.
- o New calls: FracMul, FixDiv, FracDiv, FixRound, FracSqrt, FracCos, FracSin, FixATan2, HiWord, LoWord, Long2Fix, Fix2Long, Fix2Frac, Frac2Fix, Fix2X, Frac2X, X2Fix, X2Frac.

Memory Manager

- o Optimized Purge and Compact for banks 0 and 1 and moved from RAM to ROM.
- o RAM patches and enhancements moved to ROM.
- o RAMdisk now returns bytes transferred count on DIB call.
- o SetHandleSize makes a handle temporarily un purgeable while changing handle size.

Miscellaneous Tools

RAM patches and enhancements moved to ROM:

- o AbsClamp fixes.
- o Battery RAM routines work if data bank is set to a bank other than bank data is in.
- o Firmware entry calls now return processor status in high byte instead of low byte.
- o GetAddr with ref number \$000E returns SerFlag address for SCC interrupts (useful if not using serial firmware).
- o ID manager can reuse discarded IDs.
- o Keyboard interrupts now enable VBL interrupts.
- o Munger now works with 1-char strings and returns with A=0.
- o New SysBeep call.
- o PackBytes and UnpackBytes return with A=0.
- o ReadBParam and ReadBRAM error codes corrected.
- o WriteBParam and WriteBRAM do not return error codes (this is a documentation change).
- o WriteTimeHex Bad Parameter error code is now \$31.

Monitor

- o 80-column screens maintained if break occurs and Pascal protocol in effect.
- o AppleSoft tabbing in 80-column mode now works correctly.
- o Control Panel's Maximum RAM Disk Size increased to 8128K instead of 4096K.
- o Firmware version number returned is \$1 instead of \$0.
- o Interrupts now disabled during paddle read routines.
- o Interrupts re-enabled after fatal system error (for debug DAs).
- o Mouse clamps with positive minimum and negative maximum works (e.g., \$6000 min, \$8000 max).

- o New monitor command, pound sign (#), installs monitor entry and memory peeker classic desk accessories (unless already installed), accessible via the Control Panel. Reinstalled automatically on reset; disabled by power off only.
- o New monitor command, Control-N, clears m, e, and x bits for native mode. (Control-R still switches to 8-bit, emulation mode.)
- o RESET entry point at \$00FA62 sets state register to \$0C and shadow register to \$08.
- o Shadowing of the Super Hi-Res area in Bank 1 is no longer enabled automatically.
- o WAIT routine now always exits with C=1.

QuickDraw II

RAM patches and enhancements moved to ROM:

- o 640-mode pen masks now work when portRect origin not a multiple of 8.
- o Arcs, ovals, and round rects can be drawn across bank boundaries.
- o Changes to round drawing routines: PPToPort, GetFontLore, GetROMFont, and InflateTextBuffer.
- o Current bank bytes 100...106 no longer modified by scaling and mapping calls.
- o FontFlags 1 and 2 added for pen width and color control.
- o FramePoly returns with A=0.
- o GetPort returns all four bytes of GrafPort.
- o HideCursor and ShowCursor work correctly with obscured cursor.
- o MapRgn now works on rectangular regions.
- o Pixel painting routines support QuickDraw Auxiliary Tool Set stretching and shrinking.
- o PPToPort now clips correctly to the current portRect.
- o QDStartUp and QDShutDown save and restore the scan line interrupt vector.
- o RectInRgn bug fixed.
- o ScrollRect works when the ClipRgn and VisRgn are not rectangular.
- o SetSysFont works.
- o StdPixels now returns with A=0 if the pen is not visible.
- o Text underline bug fixed.
- o TextBounds works.

New QuickDraw changes:

- o Busy flag now maintained correctly by ClosePort, OffsetRgn, InsetRgn, KillPoly, FillRect, FrameOval, PaintOval, EraseOval, InvertOval, FillOval, FrameArc, PaintArc, EraseArc, InvertArc, FillArc, FrameRRect, PaintRRect, EraseRRect, InvertRRect, and FillRRect.
- o Cursor appears in correct Super Hi-Res mode as determined by the low byte's bit 7 (320/640) of the MasterSCB.

SANE

- o Elms now can be called from any part of memory.
- o HALT exception jumping through the incorrect vector fixed.
- o Integer overflow during conversion reported.
- o STATUS call moved to ROM.

Scheduler

- o Scheduler now accepts a flush function call.
- o Task-handling RAM patch (on System Disk 1.0 and later) moved to ROM.

Serial I/O

- o First character after an XON is no longer trashed when buffering

- o is not enabled.
- o If serial mode bit 17 = 1, parity and framing error suppression are defeated.
- o Parity, baud, and data format commands work with buffering.
- o STATUS call will not report that a character is ready if the character arrives with a parity or framing error.
- o STATUS call works correctly with XON/XOFF protocol.

SmartPort

- o PR#5, following a PR#5 with I/O error (i.e., no disk in drive), now boots as expected.
- o SmartPort manipulates only Slot 6 motor on detect so the IWM can run in fast mode.

Sound

- o Fixed bug in FFStopSound call.
- o Fixed low-level RAM read/write bug.
- o Interrupts are disabled when the internal bell is active.
- o Interrupts no longer need to be disabled when accessing sound RAM.
- o New sound diagnostics with the following error codes: \$0C001 = failed RAM data test, \$0C002 = RAM address test, \$0C003 = register data test, and \$0C004 = control register test.
- o Sound Manager RAM patches and enhancements moved to ROM.

Text Tools

RAM patches moved to ROM:

- o RAM patches moved to ROM for Writing and ErrorWriting routines.
- o TextInit Illegal device error now is in 16-bit mode instead of 8.

Tool Locator

- o Optimized tool dispatcher.
- o ROM tools present on a memory expansion card are installed.

Changes from ROM 1 to ROM 3

ROM 3 is 256K (double the size of ROM 1) and contains several tools which do not exist in ROM 1. The patch file TS3 fixes known bugs in ROM 3 which were discovered after it was frozen. ROM 3 tools are basically System Software 5.0 tools, and the System Software 5.0 documentation covers these tools in detail. This Note only documents non-tool changes.

AppleDisk 3.5 and SmartPort

- o Use new routines for all block reads to fast RAM to eliminate double buffering.
- o The extended DIB status call returns the device subtype byte \$C1.
- o Fixed anomalies described in SmartPort Technical Note #6, Apple IIGS SmartPort Errata.
- o Fixed a ROM 1 bug that caused Write Protected to be returned with higher priority than Device Offline for the ProDOS STATUS call.

AppleTalk

- o AppleTalk moved to slots 1 and 2 from slot 7.

Control Panel CDA

- o The original Options menu is now the Keyboard menu and does not contain mouse parameters.
- o A new Mouse menu is present. The new keyboard microcontroller allows finer control of mouse tracking, so a selection procedure better than yes or no is present. Parameters are also available to set the keyboard mouse feature, which allows the numeric keypad to emulate a mouse.
- o Added an option to resize the RAM disk on the next reset in the RAM Disk menu. This option resets to No after one reboot and resizing so the RAM disk is not accidentally reformatted on every boot thereafter.
- o If slot 7 is set to AppleTalk, the Control Panel displays a warning if neither slot 1 nor slot 2 is similarly set.
- o The Printer Port and Modem Port menus now display only those parameters that may be changed if AppleTalk is the selection for those ports.
- o The RAM disk no longer has minimum and maximum settings, but rather one RAM disk size setting.

Monitor

- o Enhanced memory searching commands to automatically cross bank boundaries.
- o Added Step and Trace debugging functions.
- o Now provide vectors for the same functionality as the GS/OS System Service calls MEMORY_MOVER, DYN_SLOT_ARBITER and SET_SYS_SPEED in bank \$E1.
- o Now resize the RAM disk when the system is rebooted with the Control-Open Apple-Shift-Reset key combination.
- o Handle text page 2 shadowing and power-up bits in the new CYA chip.
- o Flash the border if the sound volume is set to zero and a beep is necessary.
- o In ROM 1 and earlier, the Miscellaneous Tools mouse firmware called the 8-bit mouse routines in the \$C400 space to do the work. In ROM 3, the 8-bit routines call the 16-bit routines to read the hardware. This change effectively means those programs which use 16-bit mouse calls (including desktop applications through the Event Manager) may use the mouse when slot 4 is set to Your Card.
- o Slots 1 and 2 may now be set to Printer, Modem, AppleTalk, or Your Card. With System Software 5.0, slot 7 does not need to be set to AppleTalk to use an AppleTalk network, although one can do it for compatibility. There is no transparent printing firmware in slot 7.
- o The Alternate Display Mode CDA no longer sets the system to fast speed when normal speed is selected in the Control Panel.
- o Added a new command, {val}=V, to set the video screen display I/O switches when resuming a program.
- o Control-T command now works as a toggle--executing it once changes to text mode, but now executing it again switches back to the previous video mode. You may change this saved video mode with the =V command.
- o Battery RAM value \$59 now controls the presence of the Visit Monitor and Memory Peeker CDAs. If this byte has the high bit set at boot time, the CDAs are automatically installed.
- o The Monitor and Memory Peeker both allow the use of Control-X to terminate a long display (i.e., a handle list or memory dump).

Serial I/O

- o XON and XOFF are no longer sent with the high bit set when buffering is enabled.
- o Terminal mode cursor is more consistent with the rest of the

system.

- o Extended Interface calls now return errors in the carry and the accumulator.

Toolbox

The following tools are now in ROM:

- o Window Manager
- o Menu Manager
- o Control Manager
- o Line Edit
- o Dialog Manager
- o Scrap Manager
- o Font Manager
- o List Manager

Further Reference

-
- o Apple IIGS Firmware Reference
 - o Apple IIGS Toolbox Reference
 - o Apple IIGS Technical Note #52, Loading and Special Memory
 - o SmartPort Technical Note #6, Apple IIGS SmartPort Errata

Apple IIGS
#27: Graphics Image File Formats

Revised by: Matt Deatherage November 1988

Written by: Steve Glass, Eagle Berns, Art Cabral,
Pete McDonald & Rilla Reynolds October 1987

This Technical Note formerly described the file formats for Apple IIGS graphics image files. File formats are now documented in Apple II File Type Notes under corresponding file types and auxiliary types:

File Type \$C0

Auxiliary Type \$0000	"PaintWorks" Packed Format
Auxiliary Type \$0001	PackBytes Packed Format
Auxiliary Type \$0002	"Apple Preferred" Packed Format

File Type \$C1

Auxiliary Type \$0000	32K unpacked picture image
Auxiliary Type \$0001	Unpacked QuickDraw II picture

Further Reference



- o Apple II File Type Notes


Apple IIGS
#28: Interface Card Design Guidelines

Revised by: Matt Deatherage
Written by: Cameron Birse

November 1988
October 1987

This Technical Note describes suggested dimensions for interface cards for the Apple IIGS and Apple IIe upgraded systems.





Apple IIGS
#29: Monochrome High-Resolution Mode

Revised by: Rilla Reynolds
Written by: Rilla Reynolds

November 1988
October 1987

This Technical Note discusses a 280 x 192 monochrome high-resolution mode available on the Apple IIGS and useful for clarifying some graphics.

You can select a 280 x 192 monochrome high-resolution mode on the Apple IIGS with the following steps:

1. Select Monochrome and 40-column from the Control Panel (which sets the 40-column soft switch and bit 5 in \$C029).
2. Select Hi-Res graphics mode (which sets GR and HIRES soft switches).
3. Read from to write to \$C05E (AN3).

To deselect the mode, read from or write to \$C05F.

A monochrome double high-resolution mode also exists on the IIGS, and you follow the same procedure outlined above to access it.

You can use the monochrome mode to display sharper graphics-mode text and fine lines for applications which do not require color. Note that Applesoft BASIC also supports the monochrome video mode.

The soft switches you must access in software to enable the monochrome high-resolution mode are as follows:

GR	\$C050
HIRES	\$C057
40COL	\$C00C (for monochrome double hi-res, use 80COL at \$C00D)
AN3 OFF	\$C05E

In addition, you must set bit 5 of the register at \$C029, and you must use a read-modify-write sequence since \$C029 is a multi-function register.

You can manipulate all of the soft switches listed above from the IIGS Monitor, except 40COL.

Apple IIgs

#30: Apple IIgs Hardware Reference Updates

Revised by: Jim Luther

September 1990

Written by: Rilla Reynolds & Jim Luther

October 1987

This Technical Note includes updates to the Apple IIgs Hardware Reference, published by Addison-Wesley. Please contact Apple II Developer Technical Support at the address listed in Apple II Technical Note #0 if you have additional corrections or suggestions for these manuals.

Changes since July 1990: Changed the description in "Signals at the Serial Ports and the Serial Communications Controller" to correctly note that the SCC can support a maximum asynchronous transmission rate of 57,600 bits per second (bps) in X16 clock mode.

There are two editions of the Apple IIgs Hardware Reference, the first edition (July 1987) which covers the original Apple IIgs only, and the second edition (1989) which covers both original Apple IIgs and the 1 MB Apple IIgs. Because page numbers have changed between the two editions and because an update to one edition may not be needed in both editions, this Note organizes corrections by chapter, always noting corrections to the Second Edition followed by corrections to the First Edition.

Chapter 3: Memory

Second Edition--Page 40, Table 3-2, Bits in the State register

First Edition--Page 36, Table 3-2, Bits in the State register

Switch the given values and descriptions for bits 7 and 2 as follows:

Bit	Value	Description
7	1	ALTZP: If this bit is 1, then bank-switched memory, stack, and direct page are in auxiliary memory.
	0	If this bit is 0, then bank-switched memory, stack, and direct page are in main memory.
2	1	LCBNK2: If this bit is 1, language-card RAM bank 2 is selected.
	0	If this bit is 0, language-card RAM bank 1 is selected.

Chapter 6: The Apple Desktop Bus

Second Edition--Page 148, after final paragraph

Add a new heading and description:

Control Panel Control Jumper

The ADB microcontroller provided with the 1 MB Apple IIgs includes an input that disables the text Control Panel (normally available via the Classic Desk Accessory menu). This feature allows the system parameters to be set and then protected from changes made via the text Control Panel. A jumper across the pins of connector S1 removes the text Control

Panel from the Classic Desk Accessory menu. All other installed classic desk accessories are still available in the Classic Desk Accessory menu when the S1 jumper is installed. The S1 connector is located near the ADB microcontroller at motherboard location F12.

Note: The S1 jumper does not prevent the system parameters from being changed with the graphic Control Panel (a new desk accessory normally available from the Apple menu of the Finder or of any other application that includes the Apple menu).

First Edition--Page 130, Table 6-9, Command byte syntax

The first row in the table should read:

x	x	x	x	0	0	0	0	Send Reset
---	---	---	---	---	---	---	---	------------

and not

A3	A2	A1	A0	0	0	0	0	Device Reset
----	----	----	----	---	---	---	---	--------------

First Edition--Page 131, Device Reset

Replace "Device Reset" with "Send Reset." The paragraph should be: "When a device receives a Send Reset command, it will clear all pending operations and data, and will initialize to the power-on state. The Send Reset command is not device-specific; it is sent to all devices simultaneously."

First Edition--Pages 138-139, Collision detection

The fourth sentence in the last paragraph should be: "By using the Listen register 3 command, the host can move the device with the activator pressed."

Chapter 7: Built-In I/O Ports and Clock

Second Edition--Page 154, Table 7-3, Disk-port soft switches

First Edition--Page 146, Table 7-3, Disk-port soft switches

\$C0E8	Drive disabled
\$C0E9	Drive enabled
\$C0EA	Drive 1 select
\$C0EB	Drive 2 select

In addition to the corrections listed for Table 7-3, the reference to "spindle motor switches" in the paragraph following the table should be replaced with "drive enable switches."

Second Edition--Page 155, Table 7-4, IWM states

First Edition--Page 146, Table 7-4, IWM states

Change the table to the following:

Q7	Q6	Drive	Operation
0	0	enabled	Read Data register
0	1	-	Read Status register
1	0	-	Read Handshake register
1	1	disabled	Write Mode register
1	1	enabled	Write Data register

1 = asserted state 0 = negated state - = do not care

First Edition--Page 146, after Table 7-4, IWM states

The following text and table should also be added:

"The drive enable switches and the drive select switches control the state of the disk port signals DR1 and DR2. The following table shows the relationship between these."

\$C0E8	Soft Switches			Disk Port Signals	
	\$C0E9	\$C0EA	\$C0EB	DR1	DR2
1	-	-	-	0	0
-	1	1	-	1	0
-	1	-	1	0	1

1 = asserted state 0 = negated state - = do not care

First Edition--Page 147, The Mode register

The IWM Mode register is a write-only register, so disregard the advice to use only a read-modify-write instruction sequence when manipulating bits.

Second Edition--Pages 156-7, Table 7-5, Bits in the Mode register

First Edition--Pages 147-8, Table 7-5, Bits in the Mode register

For Second Edition, change the description for bit 2, value 0 as shown. For First Edition, switch the given values and descriptions for bits 1, 2, and 4 as shown.

Bit	Value	Description
4	1	8-MHz read-clock speed selected.
	0	7-MHz read-clock speed selected. Set to 0 for all Apple IIgs disk accesses.
2	1	1-second timer is not selected.
	0	1-second timer selected. When the current disk drive is deselected, the drive will remain enabled for 1 second if this bit is clear.
1	1	Asynchronous handshake protocol selected; for all except 5.25-inch Apple disk drives.
	0	Synchronous handshake protocol selected; for 5.25-inch Apple disk drives.

Second Edition--Page 159, The serial ports

First Edition--Page 150, The serial ports

The first sentence should read: "The Apple IIgs has two serial ports located at the back of the computer, which may provide synchronous and asynchronous serial communications."

Second Edition--Page 160, Table 7-9, Pins on a serial-port connector

First Edition--Page 151, Table 7-8, Pins on a serial-port connector

Replace the table title and table with this table title, table and note:

Table 7-x Signal assignments for the mini 8-pin serial port connectors

Pin Number	Signal name	Signal Description
1	HSKo	Handshake output. Driven uninverted from the SCC's /DTR output. Voh = 3.6V; Vol = -3.6V; Rl = 450 ohms
2	HSKi	Handshake input or external clock. Received inverted at SCC's /CTS and /TRxC inputs. Vih = 0.2V; Vil = -0.2V; Ri = 12K ohms
3	TxD-	Transmit data (inverted). Driven inverted from SCC's TxD output; tri-stated when SCC's /RTS is not asserted. Voh = 3.6V; Vol = -3.6V; Rl = 450 ohms
4	GND	Signal ground. Connected to logic and chassis ground.
5	RxD-	Receive data (inverted). Received inverted at SCC's RxD input. Vih = 0.2V; Vil = -0.2V; Ri = 12K ohms
6	TxD+	Transmit data. Driven uninverted from SCC's TxD output; tri-stated when SCC's /RTS is not asserted. Voh = 3.6V; Vol = -3.6V; Rl = 450 ohms
7	GPI	General-purpose input. Received inverted at SCC's /DCD inputs. Vih = 0.2V; Vil = -0.2V; Ri = 12K ohms
8	RxD+	Receive data. Received uninverted at SCC's RxD input. Vih = 0.2V; Vil = -0.2V; Ri = 12K ohms

Note: Absolute values of specified voltages are minimums;
Ri is a minimum, Rl is a maximum.

Second Edition--Page 164, after Figure 7-9
First Edition--Page 155, after Figure 7-9

Add a new heading and description:

Signals at the Serial Ports and the Serial Communications Controller

The Apple IIgs has two serial ports which are compatible with most RS-232-C devices. This section describes the input and output signals provided at the serial ports. This section also discusses some input signals to the 8530 Serial Communications Controller (SCC) chip that are not described in the Apple IIgs Hardware Reference.

The transmit-data and receive-data lines of the Apple IIgs serial interface conform to the EIA standard RS-422, which differs from the more commonly used RS-232-C standard in that, whereas an RS-232-C transmitter modulates a signal with respect to a common ground, an RS-422 transmitter modulates the signal against an inverted copy of the same signal (to generate a differential signal). The RS-232-C receiver senses whether the received signal is sufficiently negative with respect to ground to be logical 1, whereas the RS-422 receiver simply senses which line is more negative than the other. An RS-422 signal is therefore more immune to noise and interference, and degrades less over distance, than an RS-232-C signal. If you ground the positive side of each RS-422 receiver and leave unconnected the positive side of each transmitter, you have essentially converted to EIA standard RS-423, which can be used to communicate with most RS-232-C devices over distances up to fifty feet, as illustrated in Figures 7-x1 and 7-x2.

.....
 . 8530 26LS32 Receivers IIGS Mini 8-pin . RS-232-C DTE Device
 . SCC & 26LS30 Drivers Serial Connector . DB-25 Connector

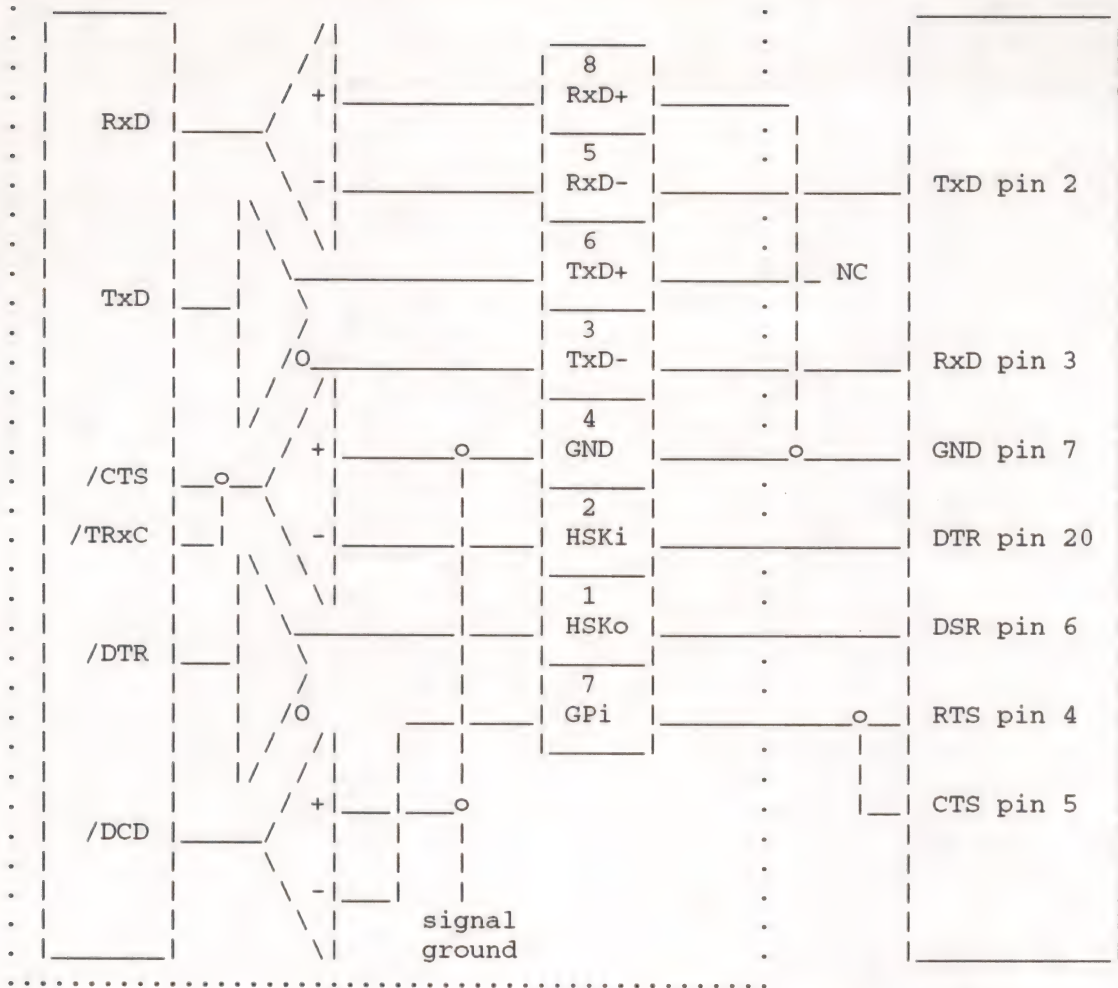
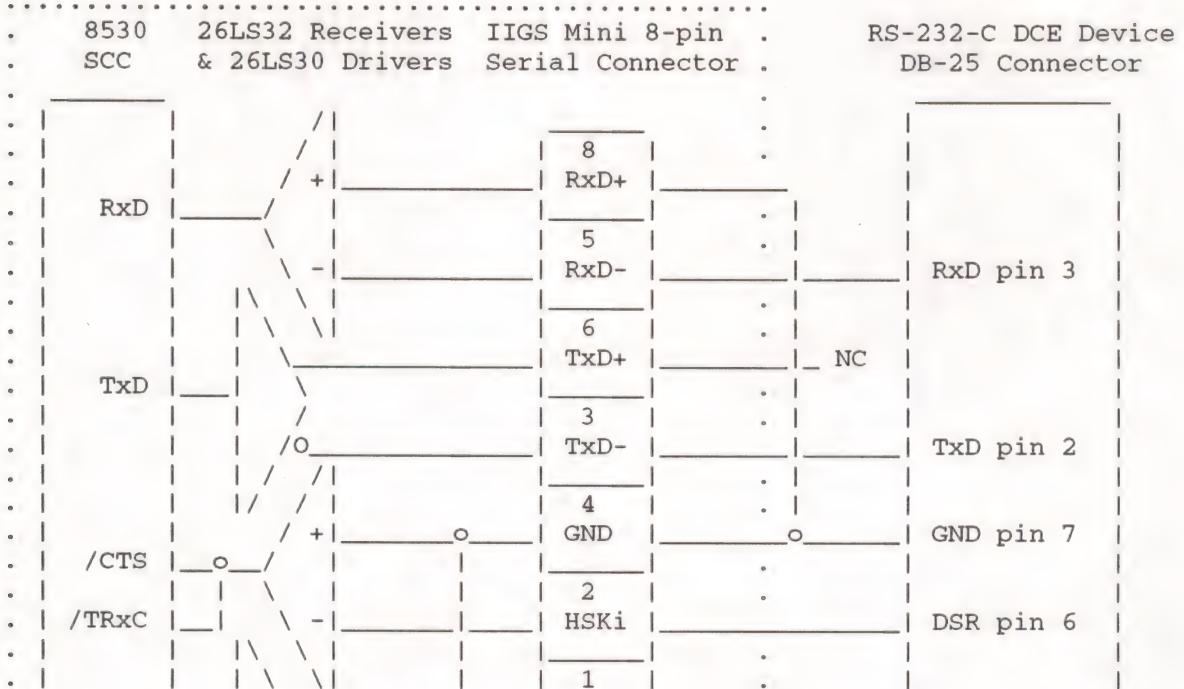


Figure 7-x1-Apple IIgs Connection to an RS-232-C DTE Device



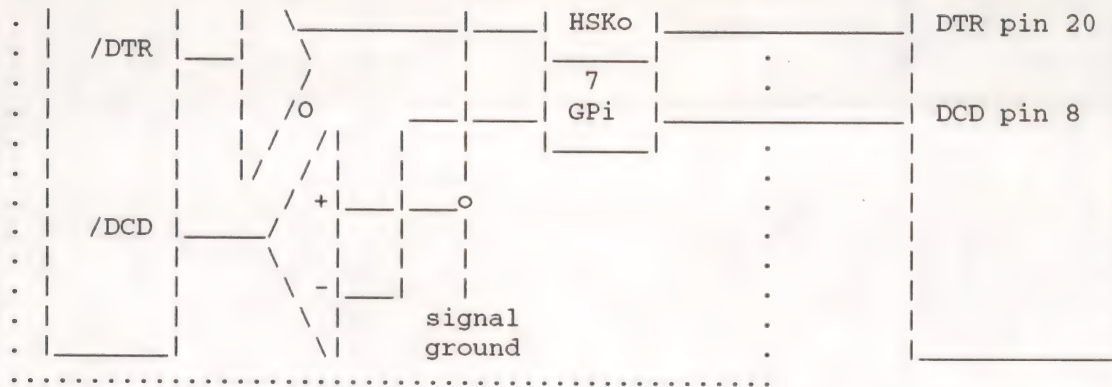


Figure 7-x2-Apple IIgs Connection to an RS-232-C DCE Device

The serial inputs and outputs of the SCC are connected to the external connectors through differential line drivers (26LS30) and receivers (26LS32). The output line drivers are tri-state devices and can be put in the high-impedance mode between transmissions to allow other devices (i.e., AppleTalk devices) to transmit over those lines. A line driver is activated by lowering the SCC's Request To Send (/RTS) output for that port.

The Handshake Output signal (HSKo, pin 1) for each Apple IIgs serial port originates at the SCC's /DTR output for that port and is driven uninverted by an RS-422 line driver (26LS30). Each port's Handshake Input signal (HSKi, pin 2) is received and inverted through a differential receiver (26LS32). The output of the differential receiver is connected to the SCC's Clear To Send (/CTS) and Transmit/Receive Clock (/TRxC) inputs for that port. HSKi is designed to accept an external device's Data Terminal Ready (DTR) handshake signal through the /CTS input. The /CTS input to the SCC can be polled by software or can be used to generate an interrupt. The HSKi line is connected to the SCC's Transmit/Receive Clock (/TRxC) input for that port, so that an external device can perform high-speed synchronous data exchange. Note that you can't use the HSKi line for receiving DTR if you're using it to receive a high-speed data clock.

Each port's General-Purpose input (GPi, pin 7) is received and inverted through a differential receiver (26LS32). The output of the differential receiver is connected to the SCC's Data Carrier Detect (/DCD) input for that port. This input can be used to provide a handshake signal from an external device to the computer. The /DCD input to the SCC can be polled by software or can be used to generate an interrupt.

Note: Because a 26LS32 differential receiver is used for the external handshake or clock signals to the SCC, the signals must be bipolar, alternating between a positive voltage and a negative voltage with respect to the internally grounded input. If a device uses ground (0 volts) as one of its handshake logic levels, the receiver interprets that level as an indeterminate state, with unpredictable results.

The SCC's Receive/Transmit Clock (/RTxC) inputs for both ports are driven by a single crystal oscillator circuit. This is accomplished by connecting a 3.6864 MHz crystal between the /RTxC and Synchronization (/SYNC) input of port A. Port B's /RTxC pin is connected to port A's /SYNC pin to drive port B's clocks from port A's oscillator circuit. Because of this single circuit, Write Register 11 (WR11) bit 7 must be set to 1 for SCC port A and must be set to 0 for SCC port B. The SCC itself is clocked at 3.58 MHz by the Apple IIgs' Color-Reference clock (CREF) at the SCC's PCLK clock input. The maximum asynchronous transmission rate supported by the SCC is 57,600

bits per second (bps) in X16 clock mode (WR4=01xxxxxx).

The SCC's Interrupt Enable In (IEI) and Interrupt Acknowledge (/INTACK) inputs are both tied to logical high in the Apple IIgs. Keeping the SCC's IEI input high enables the SCC to always generate interrupts if interrupt modes are enabled through software. Keeping the SCC's /INTACK input high leaves the SCC in Interrupt Without Acknowledge interrupt mode.

Chapter 8: I/O Expansion Slots

First Edition--Page 167, Direct memory access

DMA bank register location is \$C037.

Further Reference:

-
- o Apple IIgs Hardware Reference, both editions

Apple IIGS

#31: Redirecting Output in APW C

Revised by: Guillermo Ortiz
Written by: Guillermo Ortiz

November 1988
November 1987

This Technical Note presents a sample program which shows how to send output to different devices under the Apple Programmer's Workshop (APW) shell.

Many programmers find the ability to redirect output an especially useful feature. The following is a sample C program which allows this redirection through an APW shell command. Note that this is not applicable to MPW IIGS C since it is not part of the APW environment.

```
/*
redirect.c
Testing the shell REDIRECT command within APW C
Demonstrates how to send the output to different devices,
a disk file, the printer, and then back to the screen
last modified by Guillermo Ortiz 09/21/87

NOTE: This program checks no errors whatsoever. It expects to
be able to open the file with no problems and expects the
printer to be readily available.

Also remember that for this test to work the file has to be of
the type 'EXE' (executable from the shell only.)
*/

#include <types.h>
#include <misctool.h>
#include <stdio.h>
#include <shell.h>
#include <string.h>

char timestrg[20];          /* string to store the ascii time */
char myfile[80];           /* string to store the filename */
char str[80];              /* dummy string */
int dev=0x0001;            /* standard output */
int app=0x0000;            /* app=0 file is deleted, other will append */

PrintToFile()
{
    printf("Please enter the output filename: \n");
    gets(myfile);
    if (strlen(myfile)==0)
    {
        printf("Error in entering the filename, quit.\n");
        exit(0);
    }

    /* REDIRECT call requires pascal string */
    c2pstr(myfile);

    /* use the REDIRECT shell command to redirect the output to the file */
```

```

REDIRECT(dev, app, myfile);

/* now print a few lines of text */
printf("This is my first line of text.\n");
printf("And this is the second line.\n");
printf("Finally the third and last line of text.\n");
}

PrintToPrinter()
{
/* now redirect to output to the .PRINTER. */
REDIRECT(dev, app, "\010.PRINTER.");

printf("We should now be going to the printer.\n");
ReadAsciiTime(timestrg);
printf ("The time now is %s\n",timestrg);
}

BackToScreen()
{
/* Last REDIRECT the output back to the screen. */
REDIRECT(dev, app, "\010.CONSOLE.");

printf("The testing of REDIRECTing the output is done.\n");
ReadAsciiTime(timestrg);
printf ("The time now is %s\n",timestrg);
}

main()
{
ReadAsciiTime(timestrg);
printf ("The starting time is %s\n",timestrg);

PrintToFile();
PrintToPrinter();
BackToScreen();
}

```

Further Reference

- o Apple IIGS Programmer's Workshop Reference
- o Apple IIGS Programmer's Workshop C Reference

Apple IIGS
#32: /INH Line Anomaly

Revised by: Glenn A. Baxter & Rob Moore
Written by: Glenn A. Baxter

November 1988
December 1986

This Technical Note describes a hardware anomaly which affects the use of the /INH line on the Apple IIGS.

The Apple IIGS maps logical addresses in main and auxiliary RAM spaces to physical RAM devices in such a way that using the /INH line can cause bus contention under certain conditions. This Note describes the problem and suggests a solution strategy.

In the Apple IIGS, main memory resides within four physical 64 x 4 DRAMs. Memory is logically mapped into two separate banks of 64K x 8. The logical map of main memory is slightly different than what one might expect. Owing to the demands of new video modes on the IIGS, the DRAMs need a greater amount of time to perform their function. The easiest way to allocate time in a fixed, time-based system is to use a memory interleaving mechanism, and the IIGS implements its video in this fashion.

As a result of this interleaving scheme, the logical map of main and auxiliary memory does not correspond directly to physical DRAMs, but are split in three places. The split looks like the following:

First Physical 64K		Second Physical 64K	
Main Memory	\$0000 - \$5FFF	Auxiliary Memory	\$0000 - \$5FFF
Auxiliary Memory	\$6000 - \$9FFF	Main Memory	\$6000 - \$9FFF
Main Memory	\$A000 - \$FFFF	Auxiliary Memory	\$A000 - \$FFFF

Only part of the first physical bank of RAM is inhibited when /INH is brought low; therefore, the /INH function only works between \$0000 - \$5FFF and \$A000 - \$FFFF in main memory and \$6000 - \$6FFF in auxiliary memory. If a card attempts to inhibit main memory in the range of \$6000 - \$9FFF or auxiliary memory in the ranges \$0000 - \$5FFF or \$A000 - \$FFFF, bus contention results as both the Mega II and the 74HCT245 buffer device attempt to drive the bus simultaneously (which can damage the Mega II).

Because earlier Apple II systems do not arrange their physical memory as described above, cards which use the /INH line may be compatible with the Apple][+ and IIe, but not with the IIGS. To be compatible with all Apple II systems, a card should include an address mask that will prevent /INH from going low when the address is in the sensitive ranges of main or auxiliary memory. The following logic equation represents an appropriate blocking signal for main memory inhibition:

```
BLOCK = /A15 * A14 * A13 ;BLOCK $6000-$7FFF
        + A15 * /A14 * /A13 ;BLOCK $8000-$9FFF
```

Apple IIGS
#33: ERRORDEATH Macro

Revised by: Jim Mensch & Matt Deatherage November 1988
Written by: Allan Bell, Apple Australia & Jim Merritt December 1987

This Technical Note presents a short macro which an assembly language program can invoke to handle fatal error conditions.

Early versions of Apple-approved sample assembly language code for the Apple IIGS often invoked an APW macro named ERRORDEATH. This macro generated code that was appropriate for handling situations where program execution simply could not proceed due to "fatal" errors, such as a failure to load one or more tools that are required to display more sophisticated error dialogs or the inability to allocate sufficient direct page space for essential tool sets. The macro libraries of prototype APW systems included ERRORDEATH, but the release version does not to promote the use of more sophisticated error handling techniques in commercial software packages. The MPW IIGS release never included ERRORDEATH.

Below are two versions of ERRORDEATH; one is compatible with official standard releases of APW and the other with MPW IIGS. While Apple recommends avoiding the use of ERRORDEATH in software intended for commercial release, we feel the code is still useful for providing minimal error handling capability in prototype code and a brief, yet sophisticated, example of macro construction.

APW Assembler version:

```
MACRO
&lab      ERRORDEATH &text
&lab      bcc end&syscnt
          pha
          pea x&syscnt|-16
          pea x&syscnt
          ldx #$1503
          jsl $E10000
x&syscnt  dc il'end&syscnt-x&syscnt-1'
          dc c"&text"
          dc il'13',il'13'
          dc c'Error was $'
end&syscnt anop
MEND
```

MPW IIGS Assembler version:

```
MACRO
ErrorDeath &text
bcc @EDeathEnd
pha
pea @Message>>16
pea @Message
ldx #$1503
jsl $E10000
dc.B @EDeathEnd-@Message-1
dc.B &text
dc.B 13
dc.B 'Error Was $'
@EDeathEnd
MEnd
```

The "active ingredient" in the ERRORDEATH macro is the call to SysFailMgr (\$1503), which is made if carry is set at the time control passes to the beginning of the expanded macro code sequence. The APW and MPW IIGS assembler macro expansion mechanisms insert the value represented by the character string argument marker, &text, into the generated code stream and provide SysFailMgr with a pointer to that string. The pseudo-argument, &syscnt, generates unique labels in the positions occupied by the expressions x&syscnt and end&syscnt, which makes it possible to invoke ERRORDEATH more than once during any particular source assembly. In the MPW IIGS version of the macro, the MPW IIGS assembler creates a unique label for any label beginning with the at sign (@), effectively doing the equivalent of the &syscnt in the APW version.

To use ERRORDEATH, simply invoke it after any code sequence or subroutine call that sets the carry when it encounters an error (clears it, otherwise) and leaves an appropriate error code in the accumulator. Note that all ProDOS and Toolbox calls observe this convention. When control passes to the beginning of the ERRORDEATH code sequence, the CPU should be in full-native mode, which means the emulation bit should be clear and the accumulator and index registers should be 16-bits wide). Here is a small code segment which demonstrates invoking the macro:

```
        pushword #21          ; Dialog Manager
        pushword #0          ; Use any version
        _LoadOneTool

; If carry is now SET, following macro terminates program execution
; with the "sliding Apple" error screen.

IfWeGoofed    ERRORDEATH 'Cannot load Dialog Manager!'

; *** If no error, normal execution continues here ***
```

Apple IIGS

#34: Low-Level QuickDraw II Routines

Revised by: Dave "Evad Snoyl" Lyons, Keith Rollin,
Steven Glass, Matt Deatherage & Eric Soldan January 1991
Written by: Steven Glass May 1988

This Technical Note describes the low-level routines which QuickDraw II uses to do much of the work in standard calls and mechanisms for calling these routines and accessing their data.

Changed since November 1990: Added a Note on custom bottleneck procedures and updated information on ShieldCursor and UnShieldCursor.

QuickDraw II lets you customize low-level drawing operations by intercepting the "bottleneck procedures." QuickDraw II calls an appropriate "bottleneck proc" every time it receives a call to draw an object, measure text, or deal with pictures. For example, if an application calls PaintOval, QuickDraw II calls StdOval to do the real work, and if an application calls InvertRgn, QuickDraw II calls StdRgn to do the work.

Installing your own bottleneck procedures is a little bit tricky. The QuickDraw II SetStdProcs call accepts a pointer to a 56-byte (\$38 hex) record and fills that record with the addresses of the standard bottleneck procedures of QuickDraw II. You may modify this record by replacing those addresses with the addresses of your own custom bottleneck procedures minus one. (QuickDraw II pushes the address on the stack and executes an RTL to it, so the address in the record must point to the byte before the routine.)

Note: A custom bottleneck procedure must not begin at the first byte of a segment. If it does, then the segment could load at the beginning of a bank, and the address minus one would be in the wrong bank and RTL would transfer control to the wrong location. (See Apple IIGS Technical Note #90, 65816 Tips and Pitfalls.)

After installing your own procedures, you use SetGrafProcs to tell QuickDraw II about them. The format of this call is as follows (taken from the E16.QUICKDRAW file in APW):

ostdText	GEQU	\$00 ;	Pointer - QDProcs -
ostdLine	GEQU	\$04 ;	Pointer - QDProcs -
ostdRect	GEQU	\$08 ;	Pointer - QDProcs -
ostdRRect	GEQU	\$0C ;	Pointer - QDProcs -
ostdOval	GEQU	\$10 ;	Pointer - QDProcs -
ostdArc	GEQU	\$14 ;	Pointer - QDProcs -
ostdPoly	GEQU	\$18 ;	Pointer - QDProcs -
ostdRgn	GEQU	\$1C ;	Pointer - QDProcs -
ostdPixels	GEQU	\$20 ;	Pointer - QDProcs -
ostdComment	GEQU	\$24 ;	Pointer - QDProcs -
ostdTxBnds	GEQU	\$28 ;	Pointer - QDProcs -
ostdTxMeas	GEQU	\$2C ;	Pointer - QDProcs -
ostdGetPic	GEQU	\$30 ;	Pointer - QDProcs -
ostdPutPic	GEQU	\$34 ;	Pointer - QDProcs -

The following code fragment shows how you might replace the StdRect procedure with your own for a given window:

```

pha                                ; open a test window
pha
PushLong #MWindData                ; standard setup for NewWindow
_SetNewWindow
_SetPort

PushLong #MyProcs                  ; get a record to modify
_SetStdProcs

ldy #ostdRect                      ; get the low word of my rectangle routine
lda #myRect-1                      ; (minus one) and patch it in to the record
sta myProcs,y
lda #^myRect                        ; do the same for the high word
sta myProcs+2,y

PushLong #MyProcs                  ; install the procs
_SetGrafProcs

```

The interface to bottleneck procedures is different from the interface to other QuickDraw II routines; you do not make calls via the tool dispatcher and you pass most parameters on the direct page and in registers (rather than on the stack). To write your own bottleneck procedures, you have to know where the inputs to each call are kept and how to call the standard procedures from inside your own procedures.

The standard bottleneck procedures are accessed through vectors in bank \$E0.

StdText	\$E01E04
StdLine	\$E01E08
StdRect	\$E01E0C
StdRRect	\$E01E10
StdOval	\$E01E14
StdArc	\$E01E18
StdPoly	\$E01E1C
StdRgn	\$E01E20
StdPixels	\$E01E24
StdComment	\$E01E28
StdTxMeas	\$E01E2C
StdTxBnds	\$E01E30
StdGetPic	\$E01E34
StdPutPic	\$E01E38

When you call any of the standard procedures, the first direct page of QuickDraw II is active. If you pass variables on any direct page other than the first (direct page locations greater than \$FF), you can use a simple trick to access them. For example, to access TheFillPat (\$10E) without changing the direct page register:

```

ldx    #$100                      ;offset to second DP
lda    >$OE,X                      ;gets "DP" location $10E

```

Certain locations on the direct page are always valid:

PortRef	\$24
MaxWidth	\$20
MasterSCB	\$08
UserID	\$0A

DrawVerb is usually valid, but not always:

DrawVerb	\$38
----------	------

Each of the bottleneck procedures uses the direct page differently.

QuickDraw II has an interesting bug relating to the standard conic bottleneck procedures. If you replace any of the standard procedures with your own, QuickDraw II does not perform some of the setups it normally would before calling the standard conic procedures (stdRRect, stdOval, stdArc). For example, if you replace StdRect with a custom rectangle routine, but leave the other conic pointers alone (as shown in the code fragment above), QuickDrawII will not do all of the normal setups when calling the standard conic routines. To deal with this bug of QuickDraw II, you must patch out the additional bottleneck procedures and set up those direct pages locations yourself, or the results will not be what you expect. The QuickDraw II direct-page variables you must initialize yourself in this instance are bulleted (o) below.

StdText

DrawVerb	\$38	Describes the kind of text to draw. There are three possible values: DrawCharVerb 0 DrawTextVerb 1 DrawCStrVerb 2
TextPtr	\$DA	If the draw verb is DrawTextVerb or DrawCStrVerb, TextPtr points to the text buffer or C string to draw.
TextLength	\$D8	If the draw verb is DrawTextVerb, TextLength contains the number of bytes in the text buffer.
CharToDraw	\$D6	If the draw verb is DrawCharVerb, CharToDraw contains the character to draw.

StdLine

Y1	\$A6	Starting Y value for the line to draw
X1	\$A8	Starting X value for the line to draw
Y2	\$AA	Ending Y value for the line to draw
X2	\$AB	Ending X value for the line to draw
Rect2	\$AE	Exactly the same thing as Y1, X1, Y2 and X2 in the top, left, bottom, and right of the rectangle

StdRect

DrawVerb	\$38	One of the following five drawing verbs: Frame 0 Paint 1 Erase 2 Invert 3 Fill 4
Rect1	\$A6	The rectangle to draw in standard form (top, left, bottom, right)
TheFillPat	\$10E	The pattern to use for the rectangle if the verb is Fill

Note: The QuickDraw II Auxiliary SpecialRect call does not use the rectangle bottleneck procedures.

StdRRect

DrawVerb	\$38	One of the following five drawing verbs: Frame 0 Paint 1 Erase 2 Invert 3 Fill 4
Rect1	\$A6	The boundary rectangle for the round rectangle

OvalRect	\$295	A copy of the boundary rectangle for the round rectangle
OvalHeight	\$208	The oval height for the rounded part of the round rectangle
OvalWidth	\$20A	The oval width for the rounded part of the round rectangle
o ArcAngle	\$D2	Must be 360
o StartAngle	\$D4	Must be zero
TheFillPat	\$10E	The pattern to use for the round rectangle if the verb is Fill
StdOval		
DrawVerb	\$38	One of the following five drawing verbs:
		Frame 0
		Paint 1
		Erase 2
		Invert 3
		Fill 4
Rect1	\$A6	The boundary rectangle for the oval
OvalRect	\$295	A copy of the boundary rectangle for the oval
o OvalHeight	\$208	Must be the height of the oval
o OvalWidth	\$20A	Must be the width of the oval
o ArcAngle	\$D2	Must be 360
o StartAngle	\$D4	Must be zero
TheFillPat	\$10E	The pattern to use for the oval if the verb is Fill
StdArc		
DrawVerb	\$38	One of the following five drawing verbs:
		Frame 0
		Paint 1
		Erase 2
		Invert 3
		Fill 4
Rect1	\$A6	The boundary rectangle for the arc
o OvalWidth	\$20A	Must be the width of the boundary rectangle for the arc
ArcAngle	\$D2	The number of degrees the arc will sweep
StartAngle	\$D4	The starting position of the arc
TheFillPat	\$10E	The pattern to use for the arc if the verb is Fill
StdPoly		
DrawVerb	\$38	One of the following five drawing verbs:
		Frame 0
		Paint 1
		Erase 2
		Invert 3
		Fill 4
RgnHandleA	\$50	The handle to the polygon data structure
TheFillPat	\$10E	The pattern to use for the polygon if the verb is Fill
StdRgn		
DrawVerb	\$38	One of the following five drawing verbs:
		Frame 0
		Paint 1
		Erase 2
		Invert 3
		Fill 4
RgnHandleC	\$70	The handle to the region to draw

TheFillPat	\$10E	The pattern to use for the region if the verb is Fill															
StdPixels																	
SrcLocInfo	\$CC	The LocInfo record for the source pixel map															
DestLocInfo	\$0C	The LocInfo record for the destination pixel map															
SrcRect	\$DC	The source rectangle for the operation in local coordinates for the source pixel map (as described in the source LocInfo record)															
DestRect	\$1C	The destination rectangle for the operation in local coordinates for the destination pixel map (as described in the destination LocInfo record)															
XferMode	\$E4	The mode to use for data transfer															
RgnHandleA	\$50	The handle to the first region to which drawing is clipped (usually the ClipRgn from the GrafPort) A NIL handle is not allowed. To signify no clipping, pass a handle to the WideOpen region, which is defined as 10 bytes:															
		<table border="0"> <tbody> <tr> <td>Length</td> <td>\$A</td> <td>(word)</td> </tr> <tr> <td>-MaxInt</td> <td>-\$3FFF</td> <td>(word)</td> </tr> <tr> <td>-MaxInt</td> <td>-\$3FFF</td> <td>(word)</td> </tr> <tr> <td>+MaxInt</td> <td>+\$3FFF</td> <td>(word)</td> </tr> <tr> <td>+MaxInt</td> <td>+\$3FFF</td> <td>(word)</td> </tr> </tbody> </table>	Length	\$A	(word)	-MaxInt	-\$3FFF	(word)	-MaxInt	-\$3FFF	(word)	+MaxInt	+\$3FFF	(word)	+MaxInt	+\$3FFF	(word)
Length	\$A	(word)															
-MaxInt	-\$3FFF	(word)															
-MaxInt	-\$3FFF	(word)															
+MaxInt	+\$3FFF	(word)															
+MaxInt	+\$3FFF	(word)															
RgnHandleB	\$60	The handle to the second region to which drawing is clipped (usually the VisRgn from the GrafPort) A NIL handle is not allowed. To signify no clipping, pass a handle to the WideOpen region.															
RgnHandleC	\$70	The handle to the second region to which drawing is clipped (usually the mask region from the CopyPixels or the PaintPixels call) A NIL handle is not allowed. To signify no clipping, pass a handle to the WideOpen region.															
StdComment																	
TheKind	\$A6	The kind of input for the comment															
TheSize	\$A8	The number of bytes to put into the picture															
TheHandle	\$AA	The data to put into the picture															
StdTxMeas																	
DrawVerb	\$38	Describes the kind of text to draw. There are three possible values: <table border="0"> <tbody> <tr> <td>DrawCharVerb</td> <td>0</td> </tr> <tr> <td>DrawTextVerb</td> <td>1</td> </tr> <tr> <td>DrawCStrVerb</td> <td>2</td> </tr> </tbody> </table>	DrawCharVerb	0	DrawTextVerb	1	DrawCStrVerb	2									
DrawCharVerb	0																
DrawTextVerb	1																
DrawCStrVerb	2																
TextPtr	\$DA	If the draw verb is DrawTextVerb or DrawCStrVerb, TextPtr points to the text buffer or C string to draw.															
TextLength	\$D8	If the draw verb is DrawTextVerb, TextLength contains the number of bytes in the text buffer.															
CharToDraw	\$D6	If the draw verb is DrawCharVerb, CharToDraw contains the character to															

		measure.
TheWidth	\$DE	The resulting width should be put here.
StdTxBnds		
DrawVerb	\$38	Describes the kind of text to draw. There are three possible values: DrawCharVerb 0 DrawTextVerb 1 DrawCStrVerb 2
TextPtr	\$DA	If the draw verb is DrawTextVerb or DrawCStrVerb, TextPtr points to the text buffer or C string to draw.
TextLength	\$D8	If the draw verb is DrawTextVerb, TextLength contains the number of bytes in the text buffer.
CharToDraw	\$D6	If the draw verb is DrawCharVerb, CharToDraw contains the character to draw.
RectPtr	\$D2	Indicates the address to put the resulting rectangle.

StdGetPic

This call takes input on the stack rather than the direct page. This is the one standard bottleneck procedure which you call with the direct page register set to something other than the direct page of QuickDraw II; it is set to a part of the stack.

Stack Diagram on Entrance to StdGetPic

```

Previous Contents
DataPtr           Pointer to destination buffer
Count            Integer (unsigned) (bytes to read)
RTL Address       3 bytes
-----
Top of Stack

```

Stack Diagram just before exit from StdGetPic

```

Previous Contents
RTL Address       3 bytes
-----
Top of Stack

```

StdPutPic

This call takes input on the stack rather than the direct page; however, unlike StdGetPic, the direct page for QuickDraw II is active when you call this routine.

Stack Diagram on Entrance to StdPutPic

```

Previous Contents
DataPtr           Pointer to source buffer
Count            Integer (unsigned) (bytes to read)
RTL Address       3 bytes
-----
Top of Stack

```

Stack Diagram just before exit from StdPutPic

```

Previous Contents
RTL Address       3 bytes
-----
Top of Stack

```

Dealing with the Cursor

The cursor can get in your way when you want to draw directly to the screen. QuickDraw II has two low-level routines which help you avoid this problem:

ShieldCursor and UnshieldCursor. ShieldCursor tells QuickDraw II to hide the cursor if it intersects the MinRect and to prevent the cursor from moving until you call UnshieldCursor.

There is a bug in ShieldCursor for System Disks 4.0 and earlier. This bug is related to the routine ObscureCursor. When the cursor is obscured, ShieldCursor does not prevent the cursor from moving; therefore, the user is able to move the cursor during a QuickDraw II operation, and this movement may disturb the screen image.

Calls to ShieldCursor must be balanced by calls to UnshieldCursor. You may not call ShieldCursor successively without calling UnshieldCursor after each call to ShieldCursor. There is no error checking, so careless use of these routines will result in an unusable system.

MinRect is the smallest possible rectangle which encloses all the pixels that may be affected by a drawing call. You keep MinRect on the direct page and usually calculate it by intersecting the rectangle of the object you are drawing with the BoundsRect, PortRect, boundary box of the VisRgn, and the boundary box of the ClipRgn. You must set up MinRect yourself.

ShieldCursor also looks at two other fields on the direct page of QuickDraw II. ImageRef is a long word located at \$0E. If ImageRef does not point to \$E12000 or \$012000, QuickDraw II assumes you are not drawing to the screen, so it does not have to shield the cursor. BoundsRect is a rectangle located at \$14, and QuickDraw II uses it to translate MinRect into global coordinates. These values are generally correct, but under the following known circumstance, they are not and ShieldCursor will not function properly:

1. You have just drawn to an off-screen GrafPort with QuickDraw II.
2. You switch to a GrafPort on the screen.
3. You call ShieldCursor.

ImageRef and BoundsRect are not updated until QuickDraw II is actually committed to drawing, thus, these values are still for the off-screen GrafPort in this case, even though you switched to a GrafPort on the screen. Therefore, when you call ShieldCursor, you have to make sure that these values are current. (If these values are current, ShieldCursor will work correctly, no matter what the circumstances.)

You can find the location of the QuickDraw II direct page with the GetWAP call. For speed reasons, you may not want to make the GetWAP call for each ShieldCursor call. You may wish to get the work area pointer value after starting QuickDraw II and store it for future reference.

Calling ShieldCursor:

1. Set direct page for QuickDraw II.
2. Save the existing values of MinRect, ImageRef, and BoundsRect.
3. Set MinRect, ImageRef, and BoundsRect.
4. Let QuickDraw II know you've changed the contents of its direct page by clearing the "dirty" flags bits 14 to 0:

```
DirtyFlags    equ    $EC
                ldx    #$200                ;index to QD's third page of work
lda    DirtyFlags,x
                and    #$8000
                sta    DirtyFlags,x
```

5. JSL to ShieldCursor.
6. Restore the previous values of MinRect, ImageRef, and BoundsRect.

Note: Saving and restoring these values was not previously mentioned in this Note and in most circumstances it is not necessary. Saving and restoring is now recommended. In particular, if ShieldCursor is called inside a QuickDraw II bottleneck procedure, the system can crash if you fail to restore the contents of direct page.

Calling UnshieldCursor:

1. Set direct page for QuickDraw II.
2. JSL to UnshieldCursor.

ShieldCursor	\$E01E98
MinRect	\$00
ImageRef	\$0E
BoundsRect	\$14

UnshieldCursor	\$E01E9C
----------------	----------

Further Reference

o Apple IIGS Toolbox Reference, Volume 2

Apple IIgs
#35: Printer Driver Specifications

Revised by: Matt Deatherage September 1990
Written by: Dan Hitchens, Matt Deatherage & Suki Lee May 1988

This Technical Note describes the routines and internal structures needed to design a printer driver for the Apple IIgs system, and you should use this Note with the Apple IIgs Toolbox Reference manuals. An overview and associated parameters for each of the printer driver routines are in the Print Manager chapter, and you should refer to these for a complete picture. Changed since March 1990: Added corrections and further descriptions.

Printing Modes

There are two printing modes: immediate and deferred.

- o In immediate mode, pages are printed as they are drawn into the printing grafPort. As the application makes QuickDraw II calls, the printer driver immediately generates commands, transferring ink to page when the page is closed. This is the fastest form of printing, but only produces high-quality images on printers that can translate QuickDraw II commands to other graphic commands. For example, the LaserWriter driver translates the QuickDraw II calls into PostScript(R) calls which can produce high-quality images.
- o In deferred mode (sometimes referred to as spool mode), pages are captured to memory or disk and printed after all pages have been defined. Most printer drivers use deferred mode to create high-quality images. Since most drivers cannot obtain enough memory to image an entire page at once, they redraw page in several pieces, or bands. The printer driver creates a grafPort whose boundsRect, portRect, clipRgn, and visRgn correspond to the band and plays the picture back, thus causing the saved commands to draw only the images which fall within the band. Once the pixel image for the band is created, the printer driver converts the image to printer codes and sends the codes to the printer through the port driver.

File Structure

The user can install new printer drivers into the system by copying a printer driver file into a subdirectory called DRIVERS within the SYSTEM subdirectory. The printer driver file must be of type \$BB and have an auxiliary type of \$0001.

Print Driver Calls

A printer driver must support the following calls:

PrDefault	\$0913	Sets print record to default
PrValidate	\$0A13	Validates print record

PrStlDialog	\$0B13	Performs a style dialog
PrJobDialog	\$0C13	Performs a job dialog
PrPixelMap	\$0D13	Prints a pixel map
PrOpenDoc	\$0E13	Opens the document
PrCloseDoc	\$0F13	Closes the document
PrOpenPage	\$1013	Opens a page
PrClosePage	\$1113	Closes a page
PrPicFile	\$1213	Prints a picture file
--RESERVED--	\$1313	
PrError	\$1413	Gets the error value
PrSetError	\$1513	Sets the error value
GetDeviceName	\$1713	Gets device's name
PrDriverVer	\$2313	Gets installed driver version

Printer drivers may support the following calls if they use the new driver structure outlined below:

PrGetPrinterSpecs	\$1813	Returns printer type and characteristics
PrGetPgOrientation	\$3813	Returns page orientation

Print Driver Entry

- o For older drivers, entry is at the first byte (no offset). For newer (Print Manager 3.0 and later) drivers, the first word is \$0000, indicating a new style driver. The next word is a count of how many calls this driver supports. All drivers must support the minimum call set. Additional calls must be supported in the sequence listed (for example, if a driver supports PrGetPgOrientation, it must also support PrGetPrinterSpecs).
- o The Print Manager places an index to the correct routine in the X register (see the example and note the specific ordering of the routines).
- o There are two long return addresses (six bytes) that have been pushed onto the stack. (You must take these addresses into account to access the parameters and to return correctly.)

Example

```

StartOfNewDriver    START

                   dc i2'0'                ; new style driver
                   dc i2'(ListEnd-PrDriverList)/4' ; count

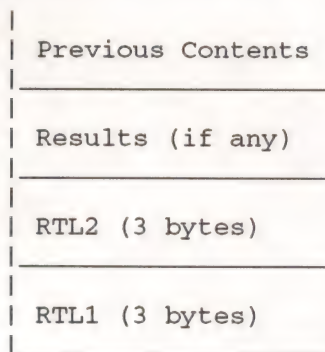
                   jmp (PrDriverList,x)

PrDriverList       dc a4'PrDefault'
                   dc a4'PrValidate'
                   dc a4'PrStlDialog'
                   dc a4'PrJobDialog'
                   dc a4'PrDriverVer'
                   dc a4'PrOpenDoc'
                   dc a4'PrCloseDoc'
                   dc a4'PrOpenPage'
                   dc a4'PrClosePage'
                   dc a4'PrPicFile'
                   dc a4'InvalidRoutine'
                   dc a4'PrError'
                   dc a4'PrSetError'
                   dc a4'GetDeviceName'
                   dc a4'PrPixelMap'
                   dc a4'PrGetPrinterSpecs'
                   dc a4'PrGetPgOrientation'

```


rtl

Figure 1 diagrams the stack just before exiting the print driver:



<-- Stack Pointer

Figure 1-Stack Prior to Exiting the Print Driver

You should do an RTL with the contents of the flags and registers set appropriately. (See the Return from Call section of the "Using The Apple Tools" chapter of the Apple IIgs Toolbox Reference.)

Print Record Structure

Since application programs often need to fiddle with parts of the print record (i.e., the values in the style subrecord), we have defined ways for applications to interpret the print record, and specifically the style subrecord.

iDev, the first word of the printer information subrecord, has two defined values for third-party printer drivers. A value of \$8001 indicates a dot-matrix printer while a value of \$8003 indicates a laser printer.

A value of \$8001 indicates that fields of the style subrecord should be interpreted as they are by the ImageWriter driver, as documented in the Apple IIgs Toolbox Reference. The first seven bits (0-6) of wDev are defined as for the ImageWriter driver. Bits 7-11 are reserved for Apple's use and must be set to zero. Bits 12-15 may be used by third-party printer drivers as necessary; these bits are set to zero in Apple's drivers.

A value of \$8003 indicates that fields of the style subrecord should be interpreted as they are by the LaserWriter driver. The first four bits (0-3) of wDev are defined as for the LaserWriter driver. Bits 4-11 are reserved for Apple's use and must be set to zero. Bits 12-15 may be used by third-party printer drivers as necessary; these bits are set to zero in Apple's drivers.

If an application wishes to take advantages of specific features of a third-party printer driver, it has to know that it is dealing with that driver. Since all drivers look pretty much alike, the Print Manager allows you to ask for the name of the currently selected printer driver. An application may make the Print Manager call PMGetPrinterName, which is documented in Volume 3 of the Toolbox Reference. The Print Manager returns the name of the currently selected printer in a Pascal (length byte) string. The name returned is the name of the file from which the driver was loaded. If you intend to use this method to identify a driver, you must inform users not to rename the Printer Driver file on the boot disk.

For alternate driver identification, Developer Technical Support assigns new

iDev values if you feel it is absolutely necessary for your driver. Please keep in mind, however, that no application knows how to interpret style records for non-standard iDev values, and that Apple does not publish such interpretations.

Print Driver Calls

Your printer driver handles the following calls:

PrDefault (\$0913)

Description:

Fills the fields of the specified print record with default values for the printer.

Passed:

PrintRecordHandle LONG Handle to the print record

Returned:

None

Performs the following:

- o Validates that PrintRecordHandle is a handle and does nothing if not.
- o Determines the default values for the print record either through tables or calculations. The default values should take into account such things as paper size and orientation, print mode, printer type, etc.
- o Copies the default values to the print record specified by the PrintRecordHandle parameter.

PrValidate (\$0A13)

Description:

Checks the print record to see that it is valid for the currently installed printer driver.

Passed:

PrintRecordHandle LONG Handle to the print record

Returned:

ChangeFlag WORD Boolean; TRUE if the record is adjusted

Performs the following:

- o Checks to see if the print record is from this particular driver.
- o If the print record is not from this driver, it uses the default values for this driver.
- o If the print record is from this driver, it makes any changes that might be needed (i.e., style, paper size, etc.).

PrStlDialog (\$0B13)

Description:

Performs a style dialog with the user.

Passed:

PrintRecordHandle LONG Handle to the print record

Returned:

ConfirmFlag WORD Boolean; TRUE if the dialog is confirmed

Performs the following:

- o Conducts a style dialog with the user to determine the page dimensions and other information needed for page setup (the initial settings of the dialog are derived from the print record).
- o If the user confirms the dialog, the information from the dialog is saved in the specified print record, PrValidate is called, and the routine returns TRUE.
- o If the user cancels the dialog, the print record is left unchanged, and the routine returns FALSE.

Note: The following are items typically found in printer style dialogs:

- o Paper Size (US Letter, US Legal, A4 Letter, B5 Letter, International Fanfold)
- o Printing Orientation (Landscape, Portrait)
- o Vertical Sizing (Normal, Intermediate, Condensed)
- o Special Effects:
 - Font Effects (Font Substitution, Smoothing)
 - Reduction or Enlargement
 - Gaps or No Gaps between pages

Every printer style dialog should have an OK button (default) and a Cancel button.

Note: When calling other routines in your printer driver (like PrValidate), be sure to do so through the Tool Dispatcher (\$E10000 or \$E10004) so any necessary patches have an opportunity to execute.

PrJobDialog (\$0C13)

Description:

Performs a job dialog with the user.

Passed:

PrintRecordHandle LONG Handle to the print record

Returned:

ConfirmFlag WORD Boolean; True if the dialog is confirmed

Performs the following:

- o Conducts a job dialog with the user to determine the print quality, range of pages to print, and other specifications. The initial settings are derived from the previous PrJobDialog call (or initial default values) except the page range which is set to ALL, and the number of copies which is set to ONE.
- o If the user confirms the dialog, PrValidate is called, the print record is updated, and the routine returns TRUE.
- o If the user cancels the dialog, the print record is left unchanged, and the routine returns FALSE.

Note: The following are items typically found in printer job dialogs:

- o Print Quality (Best, Faster, Draft, etc.)
- o Color option
- o Pages (All, Range)
- o Copies
- o Paper Source (paper cassette, manual feed)

Every printer job dialog should have an OK button (default) and a Cancel button.

Note: When calling other routines in your printer driver (like PrValidate), be sure to do so through the Tool Dispatcher (\$E10000 or \$E10004) so any necessary patches have an opportunity to execute.

PrPixelMap (\$0D13)

Description:

Prints all or part of the specified pixel map.

Passed:

srcLocPtr	LONG	Pointer to the source LocInfo which contains the pointer to the pixel map.
srcRectPtr	LONG	Pointer to the rectangle which encloses the pixel map to be printed.
colorFlag	WORD	Boolean; FALSE if black and white, TRUE if color.

Returned:

None

Performs the following:

- o Calls DevIsItSafe (port driver call) to verify that the port is functioning and it is safe to proceed. If it is not functioning, set the internal error code to \$1302 (Port Not On) and return with an error status.
- o Saves the current grafPort.
- o Turns on the watch cursor to signal the user that it will take some time.
- o Clears the internal error code (default, if no errors occur).

You can choose to print the pixel map in any convenient fashion; one convenient way is to allocate a new print record and call your normal printing routines. This method is outlined below.

- o Gets a new handle for a print record and set it to the defaults by calling PrDefault.
- o If colorFlag is set, change the style subrecord of the print record to reflect color printing.
- o Do any initialization that might be needed by the driver.
- o Determine the intersection of the two rectangles (rectangle pointed to by srcRectPtr and the pixel map's boundary rectangle from srcLocPtr) and if there is no intersection, then nothing is to be printed.
- o Print the pixel image which is within the intersection of the two rectangles.
- o Cause a page eject to occur on the printer.
- o Do any clean up that is needed.
- o Turn off the watch cursor by calling InitCursor (or restore the previous cursor using SetCursor).
- o Restore the grafPort by calling SetPort.

PrOpenDoc (\$0E13)

Description:

This routine initializes the things needed to open a document. In deferred mode, it establishes a grafPort and makes it the current port for printing.

Passed:

PrintRecordHandle	LONG	Handle to the print record
PrinterPortPtr	LONG	Pointer to the grafPort, if desired, zero to allocate a new grafPort

Returned:

PrinterPortPtrRet

LONG

Pointer to the grafPort if the PrinterPortPtr was zero

Performs the following:

- o Calls DevIsItSafe (port driver call) to verify that the port is functioning and it is safe to proceed.
- o Turns on the watch cursor to signal the user that it will take some time.
- o Validates the print record passed by calling PrValidate.
- o Clears the internal error code (default, if nothing happens).
- o Puts up a dialog indicating that printing is occurring (or preparing to print).
- o If the user needs a grafPort, create one and internally note that one was created (PrCloseDoc needs to know that one was created here).
- o Initializes parameters (i.e., page number, document number, etc.).
- o If deferred mode, create an initial page list (an array of handles to pictures) for recording pages. You can pick an arbitrary number to start with (like 20). This assumes spooling to memory; spooling to disk will obviously be different.
- o Do other initialization that might be needed to start a print job.

Possible errors:

portNotOn \$1302 Indicates Port Not On

PrCloseDoc (\$0F13)

Description:

Closes the grafPort being used for printing. For immediate mode, this routine ends the printing job. For deferred mode, this routine ends the recording of the document to be printed.

Passed:

PrintGrafPortPtr LONG Pointer to the grafPort used for printing

Returned:

None

Performs the following:

- o Checks that the last print driver call did not cause a Port Not On error. If the error occurred, do nothing and return.
- o Call ClosePort to close the printing grafPort.
- o If the driver allocated a grafPort in PrOpenDoc, disposes of it.
- o If in immediate mode, does what is needed to shut things down.
- o Takes down the information dialog box from PrOpenDoc.

Possible errors:

portNotOn \$1302 Indicates Port Not On
prBozo \$13FF Someone unloaded the driver in the middle of the print loop

PrOpenPage (\$1013)

Description:

Begins a new page only if the page falls within the page range specified in the job subrecord.

Passed:

PrintGrafPortPtr LONG Pointer to the grafPort used for printing
PageFramePtr LONG Pointer to the scaling parameter, zero for none.

Returned:
None

Performs the following:

- o Looks at the driver's internal error value, and if an error has occurred, it returns without doing anything.
- o Increments the page number.
- o Calls SetPort to make the specified port the current port.
- o Initializes the port and zeroes the boundary rectangle so no actual drawing occurs.
- o If immediate mode, then do the following:
 - If this page is to be printed, install immediate mode procedures by doing the following:
 - o Create a procedure table (get the standard procedures from SetStdProcs).
 - o Put pointers to your procedures into the table and call the QuickDraw II routine SetGrafProcs. This causes QuickDraw II calls to call your routines instead of drawing to the pixel map associated with the grafPort.
- o If deferred mode, then do the following:
 - o If the current page is out of the page range, then return without doing anything further.
 - o If the user passes his own PageFramePtr, then get it.
 - o Open a picture by calling OpenPicture and adding its handle to the page list array described in PrOpenDoc.
 - o Set the ClipRgn and VisRgn to the sizing framing rectangle specified by PageFramePtr, or if none was specified, to the default of rPage.

Possible errors:

portNotOn	\$1302	Indicates Port Not On
prBozo	\$13FF	Someone unloaded the driver in the middle of the print loop

PrClosePage (\$1113)

Description:

This signals the end of a page.

Passed:

PrintGrafPortPtr	LONG	Pointer to the grafPort used for printing
------------------	------	---

Returned:
None

Performs the following:

- o Looks at the driver's internal error value and if a Port Not On error has occurred, it returns without doing anything.
- o If immediate mode, do the following:
 - o If the current page is within the range of pages to be printed, then cause a form feed (unless no gap was specified).
 - o If deferred mode, do the following:
 - o If there was no picture generated, then do nothing and just return.
 - o Restore the grafPort to the port saved in PrOpenPage.
 - o Do a ClosePicture to close the picture.

Possible errors:

portNotOn	\$1302	Indicates Port Not On
prBozo	\$13FF	Someone unloaded the driver in the middle of the print loop

PrPicFile (\$1213)

Description:

Prints a picture file generated in deferred mode.

Passed:

PrintRecordHandle	LONG	Handle to the print record
PrintGrafPortPtr	LONG	Pointer to the grafPort used for printing
StatusRecPtr	LONG	Pointer to the printer status record

Returned:

None

Performs the following:

- o Looks at the driver's internal error value and if a Port Not On error has occurred, it returns without doing anything.
- o If immediate mode, return without doing anything.
- o If deferred mode, then do the following:
 - o If the error code is not zero (errors) then dispose of all the recorded page images.
 - o Put up an information dialog indicating that printing is occurring.
 - o Display a watch cursor (saving the current cursor first if you like).
 - o If PrintGrafPortPtr is NIL, create one and make a note of it.
 - o Call OpenPort to make the grafPort the current port.
 - o If StatusRecPtr is NIL, use an internal one. This is to simplify your code; if the StatusRecPtr is NIL, you can reasonably choose not to use a status record at all, but this requires an extra code path.
 - o Initialize the status record and the number of copies counter.
 - o If the idle procedure pointer in the print record is NIL, point to an internal one. Again, as with the StatusRecPtr, you can choose to ignore idle procedures if no pointer is provided, but this requires an extra code path.
- o Do The Following For Each Copy:
 - o Calculate the number of bands to print one page and initialize the page counter.
 - o Do The Following For Each Page:
 - o Call the idle procedure routine and initialize the band counter.
 - o Get the handle to the picture associated with the current page.
 - o Set the dirty flag in the status record to FALSE.
 - o If manual paper feed, put up a dialog and wait for a response.
 - o Do The Following For Each Band:
 - o Call the idle procedure.
 - o Calculate the band rectangle and update iCurBand with the current band number.
 - o Call the idle procedure again.
 - o Set the imaging flag in the status record to TRUE.
 - o Call InitPort to reinitialize the port.
 - o Adjust fields in the port to cause drawing into the band buffer.
 - o Adjust fields in the location information field of the status record and calculate the sizing rectangle.
 - o Calculate the boundary rectangle for the band and set the port rectangle to it.
 - o Set the ClipRgn and the VisRgn to the sizing rectangle.
 - o Initialize the band by filling it with white space.
 - o Call DrawPicture to draw the picture into the band's rectangle.
 - o Do whatever is needed to print the pixel image in the band's rectangle.
 - o Clear the imaging flag.
 - o Calculate the next band's position.

- o Increment the band's counter and loop back if not done.
- o If a vertical gap was specified, cause a form feed.
- o Increment the page count to the next page and loop back if not done.
- o Increment the number of copies counter and loop back if not done.
- o Free any buffers that you own and close the port.
- o Dispose of the information dialog that you put up.
- o Dispose of each picture in the picture list by calling KillPicture.
- o Dispose of the picture list itself.
- o Restore the cursor.

Possible errors:

portNotOn	\$1302	Indicates Port Not On
prBozo	\$13FF	Someone unloaded the driver in the middle of the print loop

PrError (\$1413)

Description:

Gets the error code from the last Print Manager call.

Passed:

None

Returned:

LastError	WORD	Result code from last Print Manager call
-----------	------	--

Performs the following:

- o Gets the driver's internal error value (which was determined by the last driver call) and sets the return parameter LastError to it.

Possible Errors:

noError	\$0000	
PrAbort	\$0080	Indicates print job was aborted
	\$1301	Indicates missing drivers
	\$1302	Indicates Port Not On
	\$1303	Indicates No Print Record
	\$1306	Indicates PAP Connection Not Made
	\$1307	Indicates Read/Write PAP Error
	\$1308	Indicates Printer Connection Failed
prBozo	\$13FF	Someone unloaded the driver in the middle of the print loop

PrSetError (\$1513)

Description:

Sets the error value.

Passed:

ErrorNumber	WORD	Error number to be set
-------------	------	------------------------

Returned:

None

Performs the following:

- o Sets the driver's internal error value to the value of the passed ErrorNumber parameter.

GetDeviceName (\$1713)

(also known as PrChanged)

Description:

Used as a communications tool between the printer driver and port driver.

Passed:

None

Returned:

None

Performs the following:

- o Calls the port driver routine PrDevPrChanged with the printer name as input. This is necessary for drivers that work over AppleTalk. The name passed as the parameter to PrDevPrChanged should be what AppleTalk uses in an NBPLookup situation; for AppleTalk, such a name should follow Name Binding Protocol conventions.

This routine will be called by the Print Manager when your driver is first loaded so a network port driver can find devices of your type.

Applications should not make this call. When this routine will be called is not guaranteed; you can't use this as a substitute for a startup call.

PrDriverVer (\$2313)

Description:

Returns the version number of the currently installed printer driver.

Passed:

WordSpace WORD Space for results

Returned:

versionInfo WORD Printer driver's version number

Performs the following:

- o Gets the internal version number of the printer driver and returns it on the stack at versionInfo.

Note: The internal version number is stored major byte, minor byte (i.e., \$0103 represents version 1.3)

PrGetPrinterSpecs (\$1813)

Description:

Returns the type of printer and the printer's characteristics.

Passed:

WordSpace WORD Space for results
WordSpace WORD Space for results

Returned:

PrinterType WORD 0 = undefined
1 = ImageWriter or ImageWriter II
2 = ImageWriter LQ
3 = LaserWriter family
(except IIsc)
4 = Epson
\$8001 = generic dot matrix printer
\$8003 = generic laser printer
PrCharacteristics WORD Bits 15-2 = reserved, must be zero
Bits 1-0: 00 = cannot determine
01 = black and white

only
10 = color capable

Performs the following:

- o Returns characteristics intrinsic for the printer being supported.

The value returned for PrinterType should be the driver's iDev value.

PrGetPgOrientation (\$3813)

Description:

Returns the page orientation from a print record.

Passed:

WordSpace	WORD	Space for result
PrintRecordHandle	LONG	Handle to the print record

Returned:

PgOrientation	WORD	Current page orientation: 0 = portrait 1 = landscape
---------------	------	--

Performs the following:

- o Returns the page orientation from the current page setup information in the print record.

Immediate Mode Procedures

To print in the immediate mode, you need to install procedures which cause printing when you make QuickDraw II calls (as noted in PrOpenPage). This section describes the structure and parameters for these routines.

The basic idea is that your driver replaces low-level QuickDraw II routines with pointers to your own routines. For example, when someone wants QuickDraw II to draw some text (say with DrawString), QuickDraw II calls your low-level routine to draw the text. You can then print the text instead.

To install the immediate mode procedures, first create a procedure table for sixteen entries (16*4 bytes) and fill it with the standard procedures by calling SetStdProcs. Once you have the standard procedures, install the addresses of your replacement procedures into it and call SetGrafProcs. Installing your procedure addresses causes the appropriate QuickDraw II calls to call your procedures, which, in turn, perform the actual printing.

The routines that need to be written are known as QuickDraw II "bottleneck procedures." For most dot-matrix printer drivers, the one of most concern when writing immediate mode procedures is StdText. If your target device has an alternate page imaging language, you may wish to print entirely in immediate mode. In this case, you want to intercept most of the bottleneck procedures. Apple IIgs Technical Note #34, Low-Level QuickDraw II Routines, contains information on how to install these procedures. The sample code which follows shows how to replace StdPixels and StdText.

Example:

```
*****  
** Example of Immediate Mode Printer Procedures. **  
*****  
Immedprocs      Start  
  
SrcRect         equ $DC
```

```

SrcLocInfo    equ $CC
DrawVerb      equ $38
TextPtr       equ $da
TextLength    equ $d8
CharToDraw    equ $d6

```

```

;-----
;
; StdPixels Procedure (Prints Pixel maps)
;
;-----

```

```

Pixel          Entry

                phb                ;save data bank reg on stack
                phk                ;get program bank reg.
                plb                ;use as data bank reg.

                lda iPrErr         ;get errors
                beq Continue       ;branch if none
                brl ExitPixel      ;branch if errors

```

```

Continue       anop

```

```

;This gets the source rectangle and stores it at PixelRect

```

```

                ldx #6
MoveSrc        lda SrcRect,x
                sta PixelRect,x
                dex
                dex
                bpl MoveSrc

```

```

;This gets the source LocInfo and stores it at PixelLoc

```

```

                ldx #16-2
MoveLI        lda SrcLocInfo,x
                sta PixelLoc,x
                dex
                dex
                bpl MoveLI

```

```

                pushlong #PixelLoc ;push pointer to LocInfo
                pushlong #PixelRect ;push pointer to rectangle

```

```

;+++++
; Insert code here to print a pixel map
;   INPUT:   PixelLoc  LONG, Pointer to pixel LocInfo
;           PixelRect LONG, Pointer to pixels BoundsRect
;   SP->
;+++++

```

```

Exitpixel     lda #0                ;return with no errors
                clc
                plb                ;restore data bank
                rtl                ;return with long

```

```

PixelLoc      ds 16                ;pixel LocInfo
PixelRect     ds 8                 ;pixel rectangle

```

```

;-----
;
; StdText Procedure (Prints Standard Text)
;

```

```

;-----
StdText      Entry

              phb                ;save data bank reg on stack
              phk                ;get program bank reg.
              plb                ;use as data bank reg.

              pushlong #PenPos
              _GetPen            ;current pen pos. -> PenPos

;+++++
; Insert Code Here to move the printers head to the corresponding
; PenPos position (if needed).
;+++++

              pushword #0        ;space for textwidth
                                   ;(for call to _TextWidth)

              lda DrawVerb       ;get DrawVerb
              beq DoCar          ;if DrawVerb=0 then DoCar

              cmp #1
              beq Dotext2       ;if DrawVerb=1 then Dotext2
;
;We get here if it's a "C" string (DrawVerb=2)
;
DoCstring     anop
              sep #$20
              longa off
;Search down through string looking for terminator to calc. length
              ldy #0
KeepLooking   lda [TextPtr],y
              beq TheEnd
              iny
              bra KeepLooking
TheEnd        rep #$20
              longa on
              lda TextPtr+2
              pha                ;push the pointer to string
              lda Textptr
              pha
              phy                ;push the length of sting
              bra Common

;
;We get here if it's just one character (DrawVerb=0)
;
DoCar         anop
              pushword #0
              tdc
              clc
              adc #CharToDraw    ;calculate addr. of char.
              pha                ;push addr. of character
              pushword #1       ;push length of one char.
              bra Common

;
;We get here if it's a string of text (DrawVerb=1)
;
DoText2       anop
              lda TextPtr+2
              pha                ;push pointer to the string

```

```

        lda Textptr
        pha
        lda TextLength
        pha
Common   lda 5,s
        pha
        lda 5,s
        pha
        lda 5,s
        pha
;+++++
; Insert code here to print the text
;
; INPUT:  TextPointer  LONG, Pointer to text to print
;         TextLength  WORD, No. of bytes to print
; SP->
;+++++
        _TextWidth    ;get the texts width (DH)
        pushword #0   ;set (DV)=0
        _Move         ;move current pen location

ExitText  lda #0
        clc
        plb
        rtl
        ;return with no errors
        ;restore data bank
        ;return with long

PenPos   ds 4
        end
        ;pen position

```

Further Reference

-
- o Apple IIgs Toolbox Reference, Volumes 1-3
 - o Apple IIgs Technical Note #36, Port Driver Specifications
 - o Apple IIgs Technical Note #90, 65816 Tips and Pitfalls

PostScript is a registered trademark of Adobe Systems Incorporated.

Apple IIGS
#36: Port Driver Specifications

Revised by: Matt Deatherage & Suki Lee

September 1989

Written by: Dan Hitchens

May 1988

This Technical Note describes how to write your own drivers for Apple IIGS ports.

Changed since January 1989: Added description of new port driver structure.

Introduction

A port driver handles certain hardware-specific duties for the Print Manager, such as initializing firmware and handling low-level hardware handshaking protocols, if any are implemented. The port driver structure, like the printer driver structure, insulates the Print Manager from low-level details of printers and interface cards (or ports) so that the same calls work across various hardware configurations, provided drivers are installed on the boot disk.

Note that a port driver could also easily be called a card driver; the term port is used because the first ones written were for the internal ports of the Apple IIGS. A port driver could interface any printer (for which there is a printer driver) with any kind of port or peripheral card that can handle it. A familiar example would be a parallel printer interface card--a port driver for a parallel card would enable the Print Manager to print graphics to any parallel printer connected to it (provided, again, there was a printer driver for the particular printer installed).

In general, you need a port driver for each port or interface card through which you intend to print, and a printer driver for each printer to which you intend to print. On System Disk 4.0, Apple provides port driver files for the printer port (PRINTER), the modem port (MODEM), a port connected to the AppleTalk network (APPLETALK), and a parallel printer interface card (PARALLEL.CARD). Apple also provides printer drivers for the ImageWriter and ImageWriter II (IMAGEWRITER), the ImageWriter LQ (IMAGEWRITER.LQ), the LaserWriter family (LASERWRITER), and an Epson (EPSON). With this configuration, you can print to any of the printer types above through any of the ports, cards, or over AppleTalk. Other printer drivers and port drivers would extend the user's selection of available configurations.

What's in a Port Driver

File Structure

Users can install new port drivers into the system by copying a port driver file into a subdirectory called DRIVERS within the SYSTEM subdirectory or by running the Installer if the driver is supplied with a script to install it. The port driver file must be of type \$BB. There are two kinds of port drivers: local drivers, intended to drive a printer connected locally, and network drivers, which handle printers connected over an AppleTalk network. Local drivers have an auxiliary type of \$0002, and AppleTalk drivers (there should be only one) have an auxiliary type of \$0003.

Port Driver Calls

A port driver must support the following calls:

PrDevPrChanged	\$1913	
PrDevStartup	\$1A13	
PrDevShutDown	\$1B13	
PrDevOpen	\$1C13	
PrDevRead	\$1D13	
PrDevWrite	\$1E13	
PrDevClose	\$1F13	
PrDevStatus	\$2013	
PrDevAsyncRead	\$2113	(alias PrDevInitBack)
PrDevWriteBackground	\$2213	(alias PrDevFillBack)
PrPortVer	\$2413	
PrDevIsItSafe	\$3013	

Note that a network port driver has much more work to do than a regular (local) port or card driver. A local driver only has to worry about one printer, whereas a network port driver may find that there is not even a printer available on a running network. The information on network drivers is provided mostly for informational purposes; you should never find it necessary to write your own AppleTalk port driver.

Entering and Exiting a Port Driver

Entering and exiting is the same as described for the printer driver calls in Apple IIGS Technical Note #35, Printer Driver Specifications. The new driver structure described there applies as well. As of this writing, there are no optional calls a port driver may support. The documented list must be supported in its entirety.

PrDevPrChanged \$1913

Description:

The Print Manager makes this call every time the user accepts this port driver in the Choose Printer dialog.

Input: LONG printer name pointer

Direct Connect:

- o Makes sure that this port has been set up correctly in the Control Panel (parity, baud rate, etc.), and puts up an alert for the user if it has not been. Remember that if you change settings, even at the user's request, you should change the Battery RAM parameters as well, so the setting changes will be reflected when the user enters the Control Panel.

Network:

- o Copies the printer name to local storage for use in the NBPLookup function of the AppleTalk PAPopen and PAPstatus calls, usually by placing it in the AppleTalk parameter block. This function is similar to that performed by PrStartUp, except that PrDevPrChanged is called whenever the printer is changed by the user with the Choose Printer dialog.

PrDevStartUp \$1A13

Description:

This call is not required to do anything. However, if your driver needs to initialize itself by allocating memory or other setup tasks, this is the

place to do it. Network drivers should copy the printer name to a local storage area for later use.

Input: LONG printer name pointer
 LONG zone name pointer

Direct Connect:

- o Required to do nothing. This is a good place to do your own set-up tasks, if you have any.

Network:

- o Copies the printer name and the zone name to local storage for use in the NBPLookup function of the AppleTalk PAPopen and PAPstatus calls, usually by placing it in the AppleTalk parameter block.

PrDevShutDown \$1B13

Description:

This call, like PrDevStartUp, is not required to do anything. However, if your driver performs other tasks when it starts, from the normal (allocating memory) to the obscure (installing heartbeat tasks), it should undo them here. If you allocate anything when you start, you should deallocate it when you shutdown. Note that this call may be made without a balancing PrDevStartUp, so be prepared for this instance. For example, do not try to blindly deallocate a handle that your PrDevStartUp routine allocates and stores in local storage; if you have not called PrDevStartUp, there is no telling what will be in your local storage area.

Input: none

PrDevOpen \$1C13

Description:

This call basically prepares the firmware for printing. It must initialize the firmware for both input and output. Input is required so the connected printer may be polled for its status.

A network driver has considerably more work to do, including the possibility of asynchronous communications. Details are provided below.

Input: LONG completion routine pointer
 LONG reserved long

Direct Connect:

- o Initializes the firmware for input and output, preparing for reading from or writing to the printer.
- o If the completion pointer is NIL, then RTL. If it is not NIL, then perform a JSL to the completion routine.

Network:

- o Initializes the End-Of-Write parameter in the AppleTalk PAPWrite parameter block to zero. Never call AppleTalk INIT to initialize the firmware.
- o If the completion pointer is NIL, then prepares for synchronous communications. If it is not NIL, prepares for asynchronous printing.
- o Calls AppleTalk PAPopen to make connection, returning an error if one is returned to you.
- o Stores the AppleTalk Session number in the PAPRead, PAPWrite and PAPClose parameter blocks.
- o Executes an RTL if there is no completion routine (pointer is NIL), otherwise perform a JSL to the completion routine.

PrDevRead \$1D13

Description:

This call reads input from the printer.

Input: WORD space for result
 LONG buffer pointer
 WORD number of bytes to transfer

Output: WORD number of bytes transferred

Direct Connect:

- o Reads a specified number of bytes from the printer into the buffer.

Network:

- o Calls AppleTalk PAPRead to read synchronously. Since there is no completion pointer, reading from a network device must always be done synchronously. To read asynchronously, use PrDevAsyncRead.

PrDevWrite \$1E13

Description:

Writes the data in the buffer to the printer and calls the completion routine.

Input: LONG write completion pointer
 LONG buffer pointer
 WORD buffer length

Direct Connect:

- o Writes the contents of the buffer to the printer.
- o If the completion pointer is NIL, then RTL. If it is not, then perform a JSL to the completion routine.

Network:

- o If the completion pointer is NIL, then writing will occur synchronously. Otherwise, writing will occur asynchronously.
- o Calls AppleTalk PAPWrite to transfer the contents of the buffer.
- o If the completion pointer is NIL, then RTL to the caller. Otherwise, perform a JSL to the completion routine first, with the error code in the accumulator.

PrDevClose \$1F13

Description:

This call is not required to do anything. However, if you allocate any system resources with PrDevOpen, you should deallocate them at this time. As with start and shutdown, note that PrDevClose could be called without a balancing PrDevOpen (the reverse is not true), and you must be prepared for this if you try to deallocate resources which were never allocated.

Input: none

Direct Connect:

- o No required function.

Network:

- o Sets End-Of-Write parameter in AppleTalk PAPWrite parameter block to one.

PrPortVer \$2413

Description:

Returns the version number of the currently installed port driver.

Input: WORD space for result

Output: WORD Port driver's version number

Direct Connect and Network:

- o Gets the internal version number of the port driver and returns it on the stack.

Note: The internal version number is stored as a major byte and a minor byte (i.e., \$0103 represents version 1.3)

PrDevIsItSafe \$3013

Description:

This call checks to see if the port or card which your driver controls is enabled. It should check at least the corresponding bit of \$E0C02D, and checking the Battery RAM settings wouldn't hurt any either.

Input: WORD space for result

Output: WORD Boolean indicating if port is enabled

Direct Connect and Network:

- o Checks the system to see if the hardware and/or firmware for the card or port this driver controls is enabled, and returns TRUE if it is safe to proceed and FALSE if not. Note that for a port driver that controls an interface card, this call should return FALSE if the card is disabled and the port is enabled, while for a port driver which controls an Apple IIGS internal port, the returned value should be TRUE if the port is enabled and FALSE if not.

Further Reference

-
- o Apple IIGS Toolbox Reference, Volumes 1 & 2
 - o Apple IIGS Technical Note #35, Printer Driver Specifications

Apple IIGS
#37: Free-Form Synthesizer Tips

Revised by: Jim Mensch
Written by: Jim Mensch

November 1988
May 1988

This Technical Note is intended to help a person who is unfamiliar with the Apple IIGS Sound Tool Set use the Free-Form Synthesizer effectively.

The primary function of the Free-Form Synthesizer is to allow an application program to start one or more complex digitized or computed waveforms playing on the Apple IIGS without further intervention from the application. The waveform is a series of bytes, each representing the amplitude of your outgoing sound at a particular moment in time (defined by the sampling frequency you set). After a call to FFStartSound, the Sound Tool Set takes care of all chores involved in loading the DOC RAM, setting up registers, and actually playing your sound. Once playing, your sound will continue until either the Sound Tool Set encounters a NIL pointer in the waveform list, or until you call FFStopSound.

FFStartSound Parameters

FFStartSound has only two parameters: the first a Word containing channel, generator, and mode information, and the second a Pointer to a parameter block.

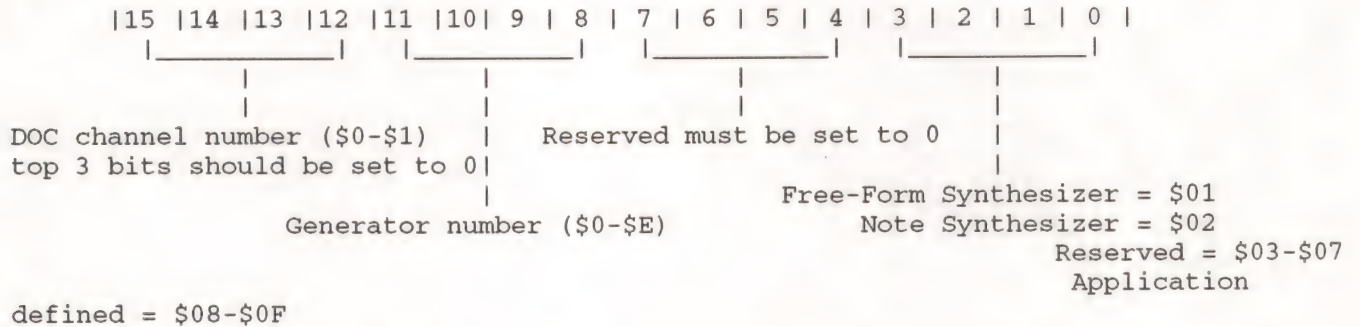


Figure 1 - Channel-Generator-Mode Word

The Channel-Generator-Mode Word is broken down into 4 nibbles. The low-order nibble specifies the particular synthesizer you are using. (Because this Note is only about the Free-Form Synthesizer, we will be using only a 1 in this nibble.) The adjacent nibble must be set to 0 for now. The next nibble specifies which generator to use. The IIGS has 15 generators from which to choose, and as the application designer, it is up to you to decide which one to use. It might be appropriate, however, to call FFGeneratorStatus first to ensure that the generator currently is available. (It could be in use already by a desk accessory or previously started sound.) The high-order nibble specifies which channel to use. The IIGS supports two separate sound channels for output. If you are using a stereo adapter, you could start up many sounds and route them to either channel 0 or channel 1 to get a full stereo effect. (The channel is ignored if you are not using a special piece of multi-channel hardware.)

The parameter block contains parameters describing the sound and how it should be played. Here is a sample Pascal definition of that parameter block:

```
FFParmBlock = record
    waveStart:Ptr;
    waveSize:Integer;
    freqOffset:Integer;
    DOCBuffer:Integer;      { High order byte significant }
    bufferSize:Integer;    { Low order byte significant }
    nextWave:^FFParmBlock;
    volSetting:Integer;
end;
```

The first parameter is a 4-byte address telling the Free-Form Synthesizer where in memory it can locate your sample data. The next parameter is a word specifying the number of 256-byte pages of sound you wish to play. The waveform data should be a series of bytes, each representing one sample. Wave tables must be exact multiples of 256 bytes.

Note: A zero value in the waveform can cause a sound to stop, so be sure to check your data to ensure that this does not happen.

The frequency offset parameter specifies the sampling frequency that the Free-Form Synthesizer should use during playback. This number can be computed by the following formula:

$$\text{freqOffset} = ((32 * \text{Sample rate in Hertz}) / 1645)$$

The frequency offset parameter is the most often misunderstood parameter, so I will explain a little about sampling rates. The sampling rate is how many samples (bytes) per second to play. If you have a digitized wave that represents 2 seconds of sound, and it takes up 44K of memory, then it was sampled at 22 kHz (which, by the way, is good for full sound reproduction). The sampling rate must be at least twice that of the maximum fundamental frequency you want to sample. However, for good sound reproduction, you may want to sample at least eight times the fundamental frequency in order to capture the higher harmonics of musical instruments and the human voice.

The DOC starting address and buffer size tell the Free-Form Synthesizer which portion of the 64K sound RAM to use as a buffer during playback. The wave is taken from your waveform in chunks and placed in sound RAM for playback. Each time the buffer nears empty, it will need to be reloaded with more sound. The size of the buffer specified determines how often the Free-Form Synthesizer must interrupt the 65816 to reload the buffer. The buffer size must be a power of two because of the way the sound General Logic Unit (GLU) specifies addresses. (The value for this parameter must also be a power of two.) A good length to use would be at least 1/10 second of sound. For example, if you were using a sampling rate of 16 kHz (16,000 samples per second), you would want a buffer at least 2,048 bytes long, or about 8 pages. It does not hurt to round this number up. You manage the DOC RAM, so you should decide what memory to use. It is usually a good idea to have multiple buffers if you have a chain of waves. (I like leaving page zero free, as the Note Synthesizer uses the data in the first 256 bytes, and accidentally placing a zero in that page could cause it to fail.)

The next wave pointer is a 4-byte pointer to the next parameter block. With this parameter you can string together many waveforms for more continuous sound, or you can make your sounds infinitely recursive by pointing back to the original wave form.

The volume setting is a word which represents the relative playback volume.

It can range from 0 to 255.

Other Tips

When you shut down the Sound Tool Set, it will stop all pending sounds, so be sure to leave ample time between starting and ending a sound. If you have a series of wave forms strung together, you can change their parameters on the fly. Changes take effect as soon as the waveform is started. (You could use this to find the correct sampling frequency of a wave, by having the next wave pointer point back to the start of your parameter block. This would cause the sound to play indefinitely. You then could change the freqOffset value, and the sound would change each time it is restarted.)

Here is a sample code segment (in APW Assembler format) that creates a 1-kHz wave in memory sampled at 16 kHz and plays it:

```
FFSound      DATA

theSound     ds      $2000          ; FFSound wave...
MyFFRecord   dc      A4'theSound'    ; address of wave
              dc      i'$20'        ; size of wave in pages..
Rate         dc      i'311'        ; 16-kHz sample rate
              dc      i'1'          ; DOC starting address
              dc      i'$0800'      ; DOC buffer size
              dc      a4'0'         ; no next wave
Voll         dc      i'$007F'      ; kinda medium..

; 1-kHz triangle wave sampled at 16 kHz one full segment
oneAngle     dc      i1'$40,$50,$60,$70,$80,$90,$A0,$B0'
              dc      i1'$C0,$B0,$A0,$90,$80,$70,$60,$50'
              End

TestFF       Start
Using        FFSound
MakeWave     ANop
              ldx     #$0000
MW0010      txa     ; get index
              and     #$000F      ; use just low nibble as index
              tay     ; into triangle wave table
              lda     oneAngle,y   ;
              sta     theSound,X   ; and store it into sound buf
              inx
              inx
              cpx     #$2000      ; we Done?
              blt     MW0010      ; nope better finish
              PushWord    #$0001
              PushLong   #MyFFRecord
              _FFStartSound
              rts
              end
```

Further Reference

o Apple IIGS Toolbox Reference, Volume 2

Apple IIgs
#38: List Controls in Dialog Boxes

Revised by: C.K. Haun September 1990
Written by: Keith Rollin, Dave Lyons & Eric Soldan May 1988

This Technical Note describes how to include a list control into a dialog box. Sample APW C source code is included.
Changes since March 1990: Changed input parameter definition for myFilterProc from long pointer to word pointer.

The need to put a list control into a dialog box is obvious. The Print Manager does it. The Font Manager does it. You may want to use one in your own application to manage a list of data base fields or spreadsheet functions. However, performing the task is not as obvious as the need.

Given the features of TaskMaster in System Software 5.0, it is now much easier to emulate a modal dialog in a normal window. If you need to add a list control to a modal dialog, you should seriously consider emulating a modal dialog with a normal window instead of using the Dialog Manager. If you use the Dialog Manager, the following procedure and sample C fragment illustrate the technique necessary for adding a list control.

Note that only one list control is allowed in a modal dialog. If you need more than one, the Dialog Manager cannot help you--create a normal window instead.

Individual Steps

Basically, there are three check-off items for putting a list control into a dialog box:

1. You must install the list explicitly into the dialog box yourself. This should be done after you have created the dialog box with a call to NewModalDialog or GetNewModalDialog. Do not install it as a UserItem or UserCtlItem. Installing it as a UserItem would cause the Dialog Manager to place an invisible custom control over the list, preventing later use of FindControl to manage it. Installing the list as a UserCtlItem does not allow the list control to be properly initialized.

Note: After you add the list control, you must not add any more dialog items.

```
InitValues()  
{  
    /* Get a Full Screen, invisible dialog window with only  
       a Quit button in it*/  
    myDialog = GetNewModalDialog(&PrintDialog);  
  
    /* Add this List Control ourselves */  
    myListHndl = CreateList(myDialog,&myList);  
  
    /* Get the handle for the Scrollbar Control */
```

```

listScrollHandle = (**myListHndl).ctlListBar;

/* Save and Zero out the RefCons */
listRefCons = GetCtlRefCon(myListHndl);
scrollRefCons = GetCtlRefCon(listScrollHandle);
ZeroRefCons(); /* This is explained below in item #3 */

/* Now show the dialog box */
ShowWindow(myDialog);
}

```

2. Because the list control is not a dialog item, a custom FilterProc must be installed for ModalDialog to test for mouse-down events. Pass the address of this routine (with the high bit set so that default handling of items is in effect) when you call ModalDialog.

```

pascal Word myFilterProc(theDialog, theEvent, theItem)
    GrafPortPtr    theDialog;
    EventRecord    *theEvent;
    word           *theItem;

{
    CtlRecHndl  tHandle;

    if ((*theEvent).what == mouseDownEvt) {
        FindControl(&tHandle, (*theEvent).where, theDialog);
        if ((tHandle == myListHndl) || (tHandle == listScrollHandle)) {

            /* Set the RefCons back to the way the list manager likes
               them */
            RestoreRefCons();
            TrackControl((*theEvent).where, (LongProcPtr) -1, tHandle);
            ZeroRefCons();

            /* Tell the Dialog Manager that we handled this event */
            return(true);
        }
        /* We didn't do anything, so return false to get Dialog Manager
           to handle this event */
        return(false);
    }
}

```

3. The Dialog Manager uses the RefCon field of its items (all of which are installed as controls). Unfortunately, the List Manager also uses the RefCon field for its own purposes. This shared use means that a judicious juggling of those values is required. This juggling is the reason for the two routines RestoreRefCons and ZeroRefCons used above.

```

/* Zero out the RefCons for the Dialog Manager */
ZeroRefCons()
{
    SetCtlRefCon(0, myListHndl);
    SetCtlRefCon(0, listScrollHandle);
}

/* Restore the RefCons for the List Manager */
RestoreRefCons()
{
    SetCtlRefCon(listRefCons, myListHndl);
}

```

```
        SetCtlRefCon(scrollRefCons,listScrollHandle);
    }
```

Note: Because the Dialog Manager currently uses the RefCon to keep track of which dialog item is identified with which particular control, zeroing the RefCon fields can cause a little confusion. Specifically, those who would like to do GetFirstDItem from within a Standard File call may get a zeroed RefCon as a result. This is true for Standard File 3.0 and later (System Software 5.0), as this is the first implementation of Standard File to use the List Manager.

Putting It All Together

Here are most of the pieces put together. InitTools and ShutDownStuff routines have been omitted, but they are straightforward.

```
char          **y,*z;
GrafPortPtr  myDialog;
ListCtlRecHndl myListHndl;
CtlRecHndl   listScrollHandle;
long         listRefCons, scrollRefCons;
```

```
#define Quit      ok
```

```
char quitStr[] = "\pQuit";
```

```
ItemTemplate quitButton = {
    Quit,
    140,450,154,590,
    buttonItem,
    quitStr,
    0,
    0,
    NULL};
```

```
DialogTemplate PrintDialog = {
    30,20,190,620,
    false,
    0,
    &quitButton,
    NULL};
```

```
char string1[] = "String1";
char string2[] = "String2";
char string3[] = "String3";
char string4[] = "String4";
char string5[] = "String5";
char string6[] = "String6";
char string7[] = "String7";
char string8[] = "String8";
```

```
MemRec myMembers[8] = {
    string1, 00,
    string2, 00,
    string3, 00,
    string4, 00,
    string5, 00,
    string6, 00,
    string7, 00,
```

```

        string8, 00);

ListRec myList = {
    40,175,102,400, /* Enclosing Rectangle */
    8,              /* Number of List Members */
    6,              /* Max Viewable members */
    3,              /* Bit Flag */
    1,              /* First member in view */
    NULL,          /* List control's handle */
    NULL,          /* Address of Custom drawing routine */
    10,            /* Height of list members */
    5,              /* Size of Member Records */
    (MemRecPtr)myMembers, /* Pointer to first element in MemRec[] */
    NULL,          /* Becomes Control's refCon */
    NULL           /* Color table for list's scroll bar */
};

/* ***** */

main()
{
    word what;

    InitTools();          /* initialize tools */
    InitValues();         /* Get dialog box. Install List control */
    do {
        what = ModalDialog((WordProcPtr)((long)myFilterProc | 0x80000000));
    } while (what != Quit);
    ShutDownStuff();
}

pascal Word myFilterProc(theDialog, theEvent, theItem)
    GrafPortPtr    theDialog;
    EventRecord     *theEvent;
    word            *theItem;

{
    CtlRecHndl      tHandle;

    if ((*theEvent).what == mouseDownEvt) {
        FindControl(&tHandle, (*theEvent).where, theDialog);
        if ((tHandle == myListHndl) || (tHandle == listScrollHandle)) {

            /* Set the RefCons back to the way the list manager
               likes them */
            RestoreRefCons();
            TrackControl((*theEvent).where, (LongProcPtr) -1, tHandle);
            ZeroRefCons();

            /* Tell the Dialog Manager that we handled this event */
            return(true);
        }
    }
    /* We didn't do anything, so return false to get Dialog Manager
       to handle this event */
    return(false);
}

/* Zero out the Refcons for the Dialog Manager */
ZeroRefCons()
{

```

```
        SetCtlRefCon(0,myListHndl);
        SetCtlRefCon(0,listScrollHandle);
    }

    /* Restore the Refcons for the List Manager */
    RestoreRefCons()
    {
        SetCtlRefCon(listRefCons,myListHndl);
        SetCtlRefCon(scrollRefCons,listScrollHandle);
    }

    InitValues()
    {
        /* Get a Full Screen, invisible dialog window with only a Quit button
           in it*/
        myDialog = GetNewModalDialog(&PrintDialog);

        /* Add this List Control ourselves */
        myListHndl = CreateList(myDialog,&myList);

        /* Get the handle for the Scrollbar Control */
        listScrollHandle = (**myListHndl).ctlListBar;

        /* Save and Zero out the RefCons */
        listRefCons = GetCtlRefCon(myListHndl);
        scrollRefCons = GetCtlRefCon(listScrollHandle);
        ZeroRefCons();

        /* Now show the dialog box */
        ShowWindow(myDialog);
    }
}
```

Apple IIGS
#39: Mega II Video Counters

Revised by: Dave Lyons
Written by: J. Rickard

July 1989
May 1988

This Technical Note describes the Mega II video output registers, which your applications can use to get information about where the beam is located on the Apple IIGS display.

Changes since November 1988: Corrected description of when VBL begins and simplified example code to read the scan line number.

The Mega II controls video timing for the Apple IIGS with a 16-bit counter split into a 7-bit horizontal and a 9-bit vertical part (Figure 1). The counter outputs are made available to programs running on the machine through two addresses in the I/O space, \$C02E for the vertical count and \$C02F for the horizontal count. These outputs can be used by a program for finer control over display update timing.




Figure 1 - Mega II Video Counter

You can see that one bit of the nine-bit vertical counter is in location \$E0C02F with the seven bits of the horizontal counter. Keep this location in mind when reading the counters.

The seven-bit horizontal counter starts at \$00 and counts from \$40 to \$7F (the sequence is \$00, \$40, \$41, ..., \$7E, \$7F, \$00, \$40, ...). The active video time consists of 40 one microsecond clock cycles starting with \$58 and ending with \$7F. Since this count changes at 980 nanosecond intervals, it will probably be of little use to most programs.

The nine-bit vertical counter ranges from \$FA through \$1FF (250 through 511) in NTSC mode (vertical line count of 262) and from \$C8 through \$1FF (200 through 511) in PAL video timing mode (vertical line count of 312). Vertical counter value \$100 corresponds to scan line zero in NTSC mode. The vertical count changes at 63.7 microsecond intervals, giving a program time to respond to a specific count before it changes. The vertical counter byte, at \$E0C02E, only changes half as often (at 127 microsecond intervals) since the lowest bit of the nine-bit counter is actually stored in the next byte (at \$E0C02F).

The nine-bit counter consists of bits VA, VB, VC, V0, V1, V2, V3, V4 and V5. Bits V0 through V5 can be read as a six-bit value. If this value is between 0 and 23, it is the line on the text screen currently being updated. Other values indicate the vertical blanking cycle is occurring. Bits VA through VC can be read as a three-bit value (0-7) indicating which scan line of a text character (characters are composed of eight lines) is currently being drawn.

The vertical counter can also be used to determine which scan line (0-191 for most video modes, including high-resolution and double high-resolution, and 0-199 for super high-resolution) is being updated at any given moment.

Example

Suppose you want to repaint a portion of the super high-resolution screen that will require more time than the vertical blanking period allows. You will have a tear in your animation when the screen's refresh cycle catches up with your drawing.

One solution to this problem would be locating the approximate place the tear occurs and starting your drawing when the system is scanning that line of graphics. Let's say you are painting an area that is about (for example) 100 pixels wide and 200 pixels tall in 320 mode, and that the tear will occur somewhere around scan line 80. To avoid the tear, you would wait until the system is scanning line 80, then you would start redrawing at the top of the screen. This way, you should be finished drawing when the system is back to scanning line 80 again and you will have flicker-free screen updating.

The tricky part is trying to determine just when the system is scanning any given scan line. One way to determine this is to examine the Mega II video counter registers at \$E0C02E (vertical) and \$E0C02F (horizontal), described above. By using some simple arithmetic you can come up with the exact scan line being updated. The following piece of code computes the current scan line number (assuming eight-bit native mode):

```
lda    >$E0C02F
asl    A                ;VA is now in the Carry flag
lda    >$E0C02E
rol    A                ;roll Carry into bit 0
```

The result (in A) is the low byte of the vertical counter. This value is 0 for the first scan line, 1 for the second scan line, etc. Values \$FA to \$FF are used twice, since you ignore the high byte of the vertical counter. (The six scan lines immediately above scan line 0 are numbered \$0FA to \$0FF, and the six above those are \$1FA to \$1FF.) The example code leaves the highest bit of the vertical counter in the Carry flag, if you really want it.

Note that the VBL interrupts always trigger at scan line 192, even in Super Hi-Res display mode, and that the \$C019 soft switch indicates vertical blanking is in effect starting at scan line 192. Be careful polling for a specific scan line number--if interrupts are enabled, it is conceivable that the system will be busy processing an interrupt every time that scan line is being scanned, so your program will hang forever waiting for it.

Setting a scan line interrupt is another way to determine when a particular super high-resolution scan line is being drawn. However, you must be careful in turning scan line interrupts on and off so that you do not interfere with the cursor in QuickDraw II (which uses scan line interrupts).

Further Reference

-
- o Apple IIGS Toolbox Reference, Volume 2
 - o Apple IIGS Technical Note #40, VBL Signal

Apple IIGS
#40: VBL Signal

Revised by: Dave Lyons
Written by: Rob Moore & Rilla Reynolds

July 1989
May 1988

This Technical Note discusses reading the VBL signal to accomplish smooth animation.

Changes since November 1988: Noted that vertical blanking does not begin when you might expect on the Apple IIGS and removed references to the Apple IIc.

Applications can accomplish smooth animation on the Apple IIGS and Apple IIe by changing the data on the screen during the time the system is tracing the unusable area of the display. This time is called "vertical blanking" or "VBL" in this Note. You can determine the state of the VBL signal by reading location \$C019.

On the Apple IIGS, the \$C019 sense of the VBL signal differs from the IIe. On the IIGS, the screen is blanked when the most significant bit of \$C019 is high (greater than 127 or \$7F), while on the IIe, the screen is blanked when the bit is low (less than 128 or \$80).

A VBL interrupt also is available on Apple II systems via the Apple IIGS Miscellaneous Tool Set or mouse firmware, the Apple IIe mouse card, and the Apple IIc mouse firmware.

On the Apple IIGS, vertical blanking begins at scan line 192 regardless of the display mode. When the Super Hi-Res display is visible, vertical blanking begins eight scan lines before the bottom of the display area. If the VBL interrupt is enabled, it triggers at scan line 192.

Further Reference

- o Apple IIGS Technical Note #39, Mega II Video Counters

Apple IIGS

#41: Font Family Numbers

Revised by: Matt Deatherage & Keith Rollin
Written by: Rilla Reynolds & Jeff Erickson

November 1990
May 1988

This Technical Note lists fonts and font family numbers as well as considerations when printing to a LaserWriter printer and a word of caution about using font family numbers.

Changes since November 1988: Added information about the font family numbering convention used by those who assign font family numbers.

The following table lists fonts and their corresponding font family numbers. All family numbers are listed in decimal format except the first three.

ID	Family Name	ID	Family Name
\$FFFD	Chicago	12	Los Angeles
\$FFFE	Shaston	13	Zapf Dingbats*
\$FFFF	(no font)	14	Bookman*
0	System Font	15	Helvetica Narrow*
1	System Font	16	Palatino*
2	New York	18	Zapf Chancery*
3	Geneva	20	Times*
4	Monaco	21	Helvetica*
5	Venice	22	Courier*
6	London	23	Symbol*
7	Athens	24	Taliesin
8	San Francisco	33	Avant Garde*
9	Toronto	34	New Century Schoolbook*
11	Cairo		

Fonts denoted with an asterisk (*) are resident in the ROM on the LaserWriter Plus, IINT and IINTX printers. The name of Times on these printers is actually Times-Roman. The decimal font family ID for Shaston (a modified Helvetica) is 65534 (-2), not 65524 as documented in the Font Manager chapter of the Apple IIGS Toolbox Reference.

When printing to a LaserWriter printer with the font substitution option turned on, the system substitutes Times, Helvetica, and Courier for thescreen fonts New York, Geneva, and Monaco respectively.

Prior to System Software 3.2, all non-LaserWriter fonts (except New York, Geneva, and Shaston) were converted to Courier when printing. With System Software 3.2 and later, the LaserWriter driver prints bitmap versions of the screen fonts if they are non-LaserWriter fonts unless it is driving an original LaserWriter printer. In this case, fonts which are in ROM on later LaserWriter printers are converted to Courier unless you download a PostScript version of the font prior to printing. This difference is a limitation of the current LaserWriter driver and it occurs even if the font substitution option is turned off. With System Software 5.0 and later, the LaserWriter driver uses fonts previously downloaded, although it does not download PostScript fonts itself.

Font Family Number Conventions

By convention, font family numbers that have the high bit set are designed for the 5:12 aspect ratio of the Apple IIgs computer. Font family numbers with the high bit clear are designed for computers with a 1:1 pixel aspect ratio, such as the Macintosh. Fonts designed for a 1:1 pixel aspect ratio appear "tall and skinny" when displayed on an Apple IIgs.

Some third-party font packages were released before this convention was defined; therefore, font family numbers between 1000 and 1200 (decimal) do not adhere to this convention.

Caution

Font family numbers can be arbitrary numbers which the system assigns to fonts. We recommend that you always ask for a font by name (with the Font Manager call GetFamNum), then use the returned family number as input to those calls which require it. (On the Macintosh, the Font/DA Mover checks to see if a font family number is already in use by the system when it installs fonts. If it finds that a number is already in use, it changes the current font number to an unused number. If you move a font from the Macintosh to the IIgs, the font family number is likely to be arbitrary, as is the font family number of any user-created fonts.

Further Reference

-
- o Apple IIgs Toolbox Reference, Volumes 1 & 2

Apple IIGS
#42: Custom Windows

Written by: Dan Oliver & Keith Rollin

November 1988

This Technical Note describes custom windows which are now supported with Window Manager version 2.2. This Note supersedes all prior documentation on custom windows.

With Window Manager version 2.2 or later, which is available on Apple IIGS System Disk 3.2 and later, you may now define your own type of window or window shape, such as a round or hexagonal window. You also may define a window which performs tasks that would normally be handled by an application.

To define your own type of window, a custom window, you must write a routine that performs some window functions. This routine is a window definition procedure (defProc), and in this case it is a custom window defProc. When the Window Manager needs to do something window specific, it calls your defProc.

The window defProc is a good part of the Window Manager, and writing one is not an easy task. A window defProc must perform complicated tasks that are very dependent on the state of the machine, and it must be very careful not to disturb the state of the machine. One of the problems in writing a defProc is knowing when it can do something and when it cannot. It is almost impossible to document all of the combinations of calls that you can or cannot make from one part or another of the defProc, and even if all cases were found, the resulting document would read like something from an obscure government bureau and probably be even harder to understand.

Now that you know writing a defProc is tough, here's how to make things as easy as possible. Try to understand how the system interacts with the defProc and work with the system. For example, a defProc is called to hit test window parts when the user presses the mouse button. The Window Manager will pass that part back to the defProc to perform drawing while the Window Manager is tracking the pressed button. The defProc could keep control when asked to hit test and perform the tracking itself, but since this is not how the system is designed to work, your defProc will be hard to write, may not ever work correctly, and may break in future versions of the Window Manager. Try to stay on the path outlined in this Technical Note. Also understand that the interface to definition procedures is as general as possible to allow them to perform tasks which are as yet unknown. To allow for this future growth, the outlined path is not always a clear path.

Another way to make things easier is to write conservative code. Do not assume things like the data bank being set to something nice when the defProc is called or the caller restoring the direct page pointer upon return if you have changed it. Use caution. A defProc can be very difficult to debug because it is not very linear and can be called when you least expect.

Interaction Between the Window Manager and TaskMaster

The Window Manager and TaskMaster actually do much less than many people think since window definition procedures perform most of the tasks. The definition procedures handle such things as title bars, information bars, and scroll

bars, while the Window Manager and TaskMaster support these things by passing requests to the defProc in standard ways. The Window Manager knows that windows have some shape, overlap, may contain parts, may be invisible, and are created and deleted, but it does not know much else. TaskMaster knows to call GetNextEvent and performs some tasks, but much of what many people consider TaskMaster is contained in the standard document window defProc. In addition to the list mentioned above, the defProc handles calling TrackGoAway and scrolling the content. The remainder of this Note describes what is expected of a defProc and when.

Telling the Window Manager About Your Window

You tell the Window Manager about your custom window when NewWindow creates it. Instead of passing the parameter list defined in NewWindow, you pass a pointer to a custom window parameter list. A custom window parameter list is defined as follows:

paramID	WORD	ID of parameter list, zero for custom.
newDefProc	LONG	Address of your custom defProc.
newData	BYTE[n]	Additional data defined by your defProc.

NewWindow checks the paramID field and calls your defProc with the pointer to the parameter list. See the wNew operation under Calling the Custom DefProc for more information.

Once NewWindow creates the window, the Window Manager will always know that it is defined by your defProc.

Calling the Custom defProc

A window defProc is called with the following items on the stack:

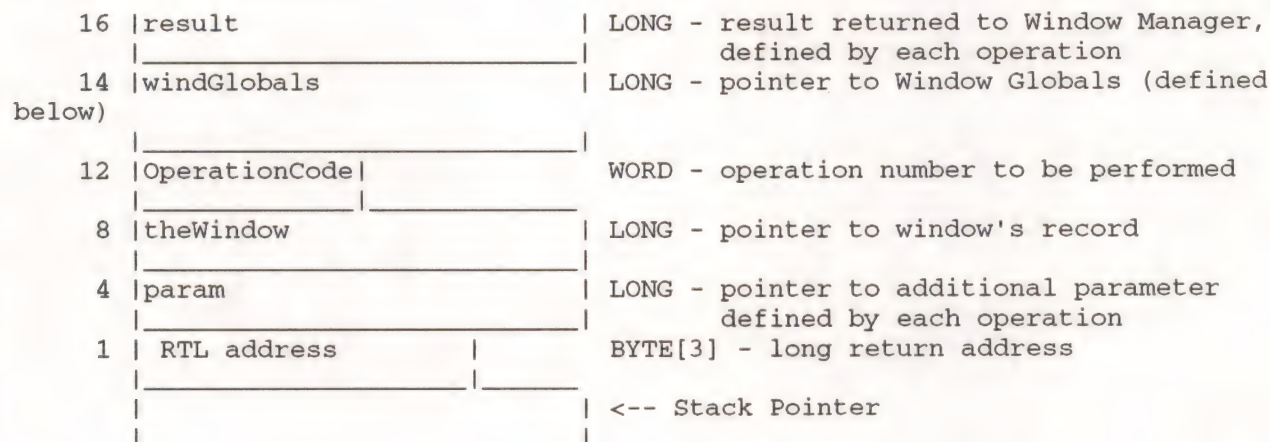


Figure 1 - Stack Prior to Calling a Window defProc

The defProc must return with the carry flag clear if there was no error or with the carry flag set and the y register set with an error code if there was an error.

Window globals (windGlobals) is a pointer to a table of variables which the Window Manager maintains for use by the defProc. The table is defined as follows:

lineW	WORD	Width of vertical lines (size depends on video mode).
titleHeight	WORD	Height of a standard title bar.

titleYPos	WORD	Y offset for the title (in system font) to center in a standard title bar.
closeHeight	WORD	Height of the close box icon.
closeWidth	WORD	Width of the close box icon.
defWindClr	LONG	Pointer to the default window color table.
windIconFont	LONG	Handle of the current window icon font.
screenMode	WORD	TRUE if 640 mode, FALSE if 320 mode.
pattern	BYTE[32]	Temporary pattern buffer.
callerDpage	WORD	Direct page pointer of the last caller to TaskMaster.
callerDataB	WORD	Data bank of the last caller to TaskMaster (bank in both bytes).

Operation numbers are as follows (each operation is described later in its own section):

wDraw	0	Draw the window's frame.
wHit	1	Tell in what region the mouse button was pressed.
wCalcRgn	2	Calculate wStrucRgn and wContRgn.
wNew	3	Complete the creation of a window.
wDispose	4	Complete the disposal of a window.
wGetDrag	5	Return address that will draw the outline of the window while dragging.
wGrowFrame	6	Draw the outline of a window being resized.
wRecSize	7	Return size of the additional space needed in the window record.
wPosition	8	Return RECT that is the window's portRect.
wBehind	9	Return where the window should be placed in the window list.
wCallDefProc	10	Generic call to a defProc, defined by the defProc.

wDraw, Operation 0

The wDraw operation draws the window's frame and is only called for visible windows. This operation draws in local coordinates in the current GrafPort, which is the Window Manager's GrafPort. When the drawing is finished, the only states of the GrafPort that may have changed are the pen pattern, the fill pattern, and the pen size, as all other states must be the same as when the defProc was called. This means that if you change the font to print some text, you must save and restore the original font. For the pen, PenNormal will restore the pen to an acceptable state.

Param is defined as follows:

Bit 31	1 to highlight the indicated part, 0 to unhighlight.
Bits 0-30	The part to draw (either highlighted or unhighlighted):
0	Draw the window's entire frame, including any frame controls and the items listed below. Note that you should check the window's fHilited flag to determine how to draw the frame.
1	Draw the go-away region.
2	Draw the zoom region.
3	Draw the information bar.

Result returned must be zero and the carry flag must be clear.

The Window Manager will draw the content.

Need to Redraw Your Window?

If your custom window defProc gets called to change some item in its window

record (see wCallDefProc below), you may want to redraw your window. For instance, if your application makes a SetWTitle call, you would want to draw the name of the new title on the screen.

The routine wCallDefProc can call the wDraw routine to do this drawing. However, it should bracket the calls to wDraw with two Window Manager calls that save and restore some internal variables:

StartFrameDrawing	\$5A0E
PUSH:LONG	Pointer to the window record (not the GrafPort)

This call does the setup for drawing a window frame and is only called by a window definition procedure before drawing the frame. You should call EndFrameDrawing when finished drawing.

EndFrameDrawing	\$5B0E
No input or output	

This call restores the Window Manager variables after a call to StartFrameDrawing and is only called by a window definition procedure after drawing a window frame.

wHit, Operation 1

The wHit operation is called to hit test the window's frame. Given a set of screen coordinates, this operation should return what part, if any, of the window is at that coordinate. This operation is only called for visible windows. The current port will be that of the Window Manager and the window frame will be in local coordinates.

Param is defined as:

Bits 0-15	Vertical (Y) coordinate in local coordinates.
Bits 16-31	Horizontal (X) coordinate in local coordinates.

Result returned must be one of the following values and the carry flag must be clear:

wNoHit	0	Not on the window at all.
wInDrag	20	Coordinates are in the window's drag region (title bar).
wInGrow	21	Coordinates are in the window's grow region (size box).
wInGoAway	22	Coordinates are in the window's go-away region (close box).
wInZoom	23	Coordinates are in the window's zoom region (zoom box).
wInInfo	24	Coordinates are in the window's information bar.
wInFrame	27	Coordinates are in the window, but not in any of the other areas.
xx		Any code the application can handle (bit 15 is reserved for the Window Manager)

wCalcRgns, Operation 2

The wCalcRgns operation, which is called only for visible windows, is used to calculate the window's entire region (frame plus content called StrucRgn) and just its content region (called ContRgn). Both regions must be set to global coordinates, and both will already be allocated with their handles stored in the window record's wStrucRgn and wContRgn fields.

Use the portRect and the boundsRect of the window's GrafPort to calculate these two regions. The port will have been set from the information passed to

NewWindow along with any size changes. A method for obtaining the global RECT of the content is given below. Refer to the QuickDraw II chapter in the Apple IIGS Toolbox Reference for a full description of ports. When calculating the regions, do not change the clip region (ClipRgn) or the visible region (VisRgn) of the GrafPort.

Param is not defined and should not be used.

Result returned must be zero and the carry flag must be clear.

IN: window = pointer to window record.
OUT: rect = global RECT of window's content.

```
    ldy    #wPort+portRect+y1
    lda    [<window],y
    ldy    #wPort+portInfo+boundsRect+y1
    sec
    sbc    [<window],y
    sta    <rect+y1
;
    ldy    #wPort+portRect+x1
    lda    [<window],y
    ldy    #wPort+portInfo+boundsRect+x1
    sec
    sbc    [<window],y
    sta    <rect+x1
;
    ldy    #wPort+portRect+y2
    lda    [<window],y
    ldy    #wPort+portInfo+boundsRect+y1
    sec
    sbc    [<window],y
    sta    <rect+y2
;
    ldy    #wPort+portRect+x2
    lda    [<window],y
    ldy    #wPort+portInfo+boundsRect+x1
    sec
    sbc    [<window],y
    sta    <rect+x2
```

Although there are other ways to obtain the global RECT of the content, this example gives the correct method. You should never rely on the top and left side of the portRect being zero.

wNew, Operation 3

The wNew operation is called to perform any additional initialization that may be required for a custom window. The following items are already done for the window:

- o If a window record is supposed to be allocated, it is. All fields, other than those fields listed below, are set to zero
- o A port opens in the window record's wPort field.
- o The window is added to the Window Manager's window list, and the wNext field is set.
- o The wDefProc, wStrucRgn, wContRgn and wUpdate regions are set with the handles of the allocated regions. It is the responsibility of the defProc to define the shape of the wStrucRgn and wContRgn regions.
- o The fAllocated and fHilited bits in the wFrame field of the window record are set (see the window record definition for a definition of these bits)

and should not be disturbed; all other bits in wFrame are set to zero. The defProc should set the fCtlTie, fVis and fQContent bits, and it can set and use other bits in the wFrame field as it wishes.

- o It is the responsibility of the defProc to set the wRefCon, wContDraw, and wFrameCtrls fields, the bits already mentioned in the wFrame field, and any other fields which it defines in the wCustom part of the window record.

Param is a pointer to the parameter list pointer which was passed to NewWindow.

Result returned must be zero and the carry flag must be clear.

wDispose, Operation 4

The wDispose operation is called to perform any additional disposal that may be required of a custom window. This operation is called before the Window Manager performs any disposal actions on the window.

Param is not defined and should not be used.

Result should be FALSE to continue disposal or TRUE to abort the disposal. In either case, the carry flag should be clear. Returning TRUE would be very unusual and should be carefully thought out. After returning FALSE, the Window Manager will erase the window, remove the window from the Window Manager's window list, free any controls in the window's wControls and wFrameCtl lists, free the handles in the wStrucRgn, wContRgn and wUpdateRgn fields, close the window's GrafPort, and free its record if it is allocated (see the wFrame field).

wGetDrag, Operation 5

The wGetDrag operation is called to get the address of a routine that will draw an outline of the window.

Param is not defined and should not be used.

Result returned must be the address of a frame outline routine or zero for a default frame; the default frame is the bounds RECT of the strucRgn. The frame outline routine is called from DragRect with dragRectPtr set to the bounds RECT of the strucRgn. Your routine is called with the following parameters:

```
PUSH:WORD - delta X
PUSH:WORD - delta Y
PUSH:BYTE[3] - return address
```

Your routine should draw or erase the outline of the object in its new position using the passed deltas. You have several different methods of determining whether to erase or draw and how to compute the position of the object, the easiest method being to draw the outline using XOR mode. The first time your routine is called, you draw. The next time your routine is called, you erase. Your routine should draw in the current port. The current pen pattern will be the pattern pointed to by dragPatternPtr from DragRect and the pen mode is XOR.

You also need to know where to draw the outline. One way is to offset the starting RECT (dragRectPtr) by the given deltas. You should make a copy of the bounds RECT of the strucRgn when wGetDrag is called. Modify that rectangle with the deltas to obtain the rectangle to frame.

wGrowFrame, Operation 6

The wGrowFrame operation is called to draw an outline of the window when the window is being resized.

This operation should use the current port, pen pattern, and pen mode. The frame should be drawn with only the following QuickDraw II calls: Line, LineTo, FrameRect, FrameRgn, FramePoly, FrameOval, FrameRRect, and FrameArc (the Invert equivalents to Frame could also be used). You want to use the current GrafPort setting with only certain QuickDraw II calls since this routine will be called an even number of times; the first time it is called to draw the frame and the next time to erase that which it drew the first time. If it needs to use QuickDraw II calls other than those listed above, this operation handler could keep track of odd and even calls to know whether to draw or erase the frame.

Param is a pointer to the following parameter list:

newSize	RECT	Rectangle that defines the new size.
drawFlag	WORD	TRUE to draw the frame, FALSE to erase.
startRect	RECT	Bounds of wStrucRgn when dragging started.
deltaY	WORD	Vertical movement since starting to drag (signed).
deltaX	WORD	Horizontal movement since starting to drag (signed).

Result should be:

```
.....  
| 31| 30| 29|...| 6 | 5 | 4 | 3 | 2 | 1 | 0 |  
.....  
| |  
TRUE if newSize RECT has been  +-- TRUE if frame drawn,  
recomputed, FALSE if newSize    FALSE to draw default frame.  
RECT OK.
```

The Window Manager assumes that the frame of the grow outline is the same as the bounds of the window's wStrucRgn. This RECT is stored in the startRect of the parameter list and does not change through out the dragging. The next assumption is that the window grows from the lower right corner. As the cursor moves, the lower right corner of the RECT in newSize changes. However, if these assumptions are not correct for a custom window they can be overridden by changing the RECT in newSize (by using startRect or the window's record and the deltas) and returning TRUE for bit 1 in Result. The carry flag should return clear.

wRecSize, Operation 7

The wRecSize operation is called to ask how large a window record should be allocated.

Note: The window pointer passed in theWindow is not valid for this call.

Param is the parameter list pointer that is passed to NewWindow.

Result is the number of additional bytes required in the window record. The standard window record header will always be allocated.

Example:

If your custom window needs a one word field in the window record for your own use you would return 2 in Result. The Window Manager takes Result and adds to

it the size of the standard record header of 212 bytes and allocates a window record that is 214 bytes long in this case. Your one word field is at the end of the standard window record header with an offset of 212 bytes.

If there is some error, return the carry flag set with an error code in the y register, which will cause NewWindow to abort and return the error code to the application which called it. If there is no error, return the carry flag clear.

Window Record Already Allocated?

If the window record is already allocated then Result should be the pointer to the window record with bit 31 of the pointer set to TRUE. Generally, window records are allocated (refer to Window Record Definition at the end of this Note for more information about window records).

wPosition, Operation 8

Param is the parameter list pointer that is passed to NewWindow.

Result is a pointer to the RECT that will be the window's portRect, and you should return the carry flag clear.

wBehind, Operation 9

Param is the parameter list pointer that is passed to NewWindow.

Result is where the window should be placed in the window list. A long \$FFFFFFFF means insert the window as the top window while a long \$00000000 means to insert it as the bottom window. Any other value is a pointer to the window behind which this window should be placed. You should return the carry flag clear.

wCallDefProc, Operation 10

WCallDefProc is a generic call to the defProc that is defined by the defProc. With this call a window defProc can define many special functions.

The input to the defProc is:

param = pointer to the following parameter table:

dRequest	WORD	Requested operation number.
paramID	WORD	Parameter block type: \$0000-\$7FFF reserved by system (\$0000 defined below). \$8000-\$FFFF reserved for custom defProcs.
newParam	BYTE[n]	New parameter field used by some operations.

The paramID field defines dRequest, which in turn defines newParam and the result of the wCallDefProc call. You can think of dRequest as the operation number passed to the defProc. Here is an example of how the paramID defines dRequest: if paramID is zero, dRequest 3 is defined as wSetPage (defined below); but if paramID is \$8345 (or any number other than zero), dRequest 3 could be defined as something entirely different.

The following dRequest values are defined for wCallDefProc operations with a paramID of zero. Your defProc should check for handling only these codes. In the future, codes 34 and greater may be defined, and your defProc should know not to handle them.

result = None.

Called when SetContentOrigin is called.

wSetDataSize 7
newParam = WORD - height of window's data area.
WORD - width of window's data area.
result = None.

Called when SetDataSize is called.

wSetZoomRect 8
newParam = LONG - pointer to new zoom RECT.
result = None.

Called when SetZoomRect is called.

wSetTitle 9
newParam = LONG - pointer to new title.
result = None.

Called when SetWTitle is called.

wSetColorTable 10
newParam = LONG - pointer to new color table.
result = None.

Called when setFrameColor is called.

wSetFrameFlag 11
newParam = LONG - pointer to new zoom RECT.
result = None.

Called when SetWFrame is called.

wGetOrgMask 12
newParam = None.
result = WORD - window's origin mask.

wGetMaxGrow 13
newParam = None.
result = Low word is window's maximum height when grown.
High word is window's maximum width when grown.

Called when GetMaxGrow is called.

wGetScroll 14
newParam = None.
result = Low word is number of pixels to scroll when arrow is
selected.

Called when GetScroll is called.

wGetPage 15
newParam = None.
result = Low word is pixels to scroll when page region is selected.

Called when GetPage is called.

wGetInfoRefCon 16
newParam = None.
result = Value passed to info bar draw routine.

Called when GetInfoRefCon is called.

wGetInfoDraw 17
 newParam = None.
 result = Address of info bar draw routine.

Called when GetInfoDraw is called.

wGetOrigin 18
 newParam = None.
 result = Low word is content's Y origin.
 High word is content's X origin.

Called when GetContentOrigin is called.

wGetDataSize 19
 newParam = None.
 result = Low word is window's data height.
 High word is window's data width.

Called when GetDataSize is called.

wGetZoomRect 20
 newParam = None
 result = Pointer to window's current zoom RECT.

Called when GetZoomRect is called.

wGetTitle 21
 newParam = None
 result = Pointer to window's title.

Called when SetWTitle is called.

wGetColorTable 22
 newParam = None.
 result = Pointer to window's color table.

Called when setFrameColor is called.

wGetFrameFlag 23
 newParam = None.
 result = Low word is window's wFrame field.

Called when SetWFrame is called.

wGetInfoRect 24
 newParam = LONG - pointer to place to store info bar's enclosing RECT.
 result = None.

Called when GetRectInfo is called.

wGetDrawInfo 25
 newParam = None.
 result = None.

Called when DrawInfoBar is called.

wGetStartInfoDraw 26
 newParam = LONG - pointer to place to store info bar's enclosing
 RECT.

result = None.

Called when StartInfoDrawing is called.

wGetEndInfoDraw 27

newParam = None.

result = None.

Called when EndInfoDrawing is called.

wZoomWindow 28

newParam = None.

result = None.

Called when ZoomWindow is called.

wStartDrawing 29

newParam = None.

result = None.

Called when StartDrawing is called.

wStartMove 30

newParam = WORD - new y position (global).

WORD - x position (global).

result = Low word is new y position (global).

High word is x position (global).

Called before MoveWindow moves a window.

wStartGrow 31

newParam = None.

result = None.

Called before GrowWindow tracks the growing of a window.

wNewSize 32

newParam = LONG - pointer to:

WORD - proposed new height.

WORD - proposed new width.

These two values can be changed.

result = Low word TRUE if only uncovered content should be drawn.

FALSE if entire content should be redrawn.

Called by SizeWindow before it resizes a window. The new height and width can be changed by modifying the words pointed to by the pointer in newParam.

wTask 33

newParam = LONG - pointer to task record.

WORD - result from FindWindow.

result = Low word is code returned by TaskMaster (zero if handled).

High word is task performed. Returned in TaskData if code is 0.

Called from TaskMaster when it cannot handle a task. If the user presses the mouse button over a window, TaskMaster will call FindWindow to find out what part of the window. TaskMaster will then handle the task if FindWindow returns wInMenuBar or bit 15 of the window pointer is set (system window). Otherwise, the result of FindWindow is passed to wTask to be handled or not.

If the defProc can handle the task it should do so and return zero in the low word of the result (which will be the result to the application returned from TaskMaster) and a code of the task performed in the high word of the result (which is returned to the application in its task record TaskData field). Fields in the task record may also be modified to return parameters to the application as this is the same record passed to TaskMaster.

If the defProc cannot handle the task, it should return the result from FindWindow (the second field in newParam) in the low word of the result. The high word of the result is not used.

For example, the standard document window defProc handles the following results from FindWindow if the taskMask record allows.

wInContent	Brings the window to the top.
wInDrag	Calls DragWindow.
wInGrow	Brings the window to the top. If it is already on the top, it calls GrowWindow and SizeWindow.
wInGoAway	Calls TrackGoAway.
wInZoom	Calls TrackZoom and ZoomWindow.
wInInfo	Brings the window to the top.
wInFrame	Brings the window to the top. If it is already on the top, checks if it is on one of the window's scroll bars, tracks it, and scrolls the window's content as needed.

A custom window defProc can return any code (bit 15 is used for system windows) it wants when it is called to do a hit test. This code would be that returned by FindWindow, and the application would have to know about the code if it called FindWindow instead of TaskMaster. If TaskMaster is used, the code that FindWindow returns is passed back to your defProc with a wCallDefProc and wTask. The defProc could perform any task it wanted: change colors, eject a disk, run a spelling checker, or anything else.

Window Record Definition

0	wNext	LONG - Pointer to next window record, zero is end of list.
4	wPort ///	BYTE[170] - Window's GrafPort.
174	wDefProc	LONG - Address of window's definition procedure.
178	wRefCon	LONG - Reserved for application's use.
182	wContDraw	LONG - Address of routine that will draw window's content.
186	wReserved	LONG - Reserved by the Window Manager, do not use.
190	wStrucRgn	LONG - Handle of window's structure region.
194	wContRgn	LONG - Handle of window's content region.
198	wUpdateRgn	LONG - Handle of window's update region.
202	wCtrls	LONG - Handle of first control in window's content.
206	wFrameCtrls	LONG - Handle of first control in window's frame.

210	wFrame		WORD - Flags that define window.

212	wCustom	...	BYTE[n] - Additional data space defined by
	_____	...	window's definition procedure.

The changes use some vacant space under the window port and add the wReserved field to the record for future expansion.

In addition to defining the window record, the wFrame field needs to be further defined. In the diagram below the shaded bits are reserved for use by each window defProc (the values shown are those used by the standard document window defProc). Bits not shaded are reserved by the Window Manager and are applicable to all windows.

Further Reference

- o Apple IIGS Toolbox Reference, Volume 1
- o System Disk 4.0 Release Notes

Apple IIgs

#43: Undocumented Feature of CalcMenuSize

Revised by: Matt Deatherage

March 1991

Written by: Dan Oliver

November 1988

This Technical Note documents that CalcMenuSize can accept a parameter of \$FFFF to recalculate menus with uninitialized heights and widths.

Changes since November 1988: This Note is now obsolete.

This Note formerly described how CalcMenuSize behaves when menu widths and heights are stored as \$0000 or \$FFFF. This behavior is now documented in Volume 3 of the Apple IIgs Toolbox Reference on page 37-3.

Further Reference

- o Apple IIgs Toolbox Reference, Volume 3

Apple IIgs

#44: GetPenState and SetPenState Record Error

Revised by: Matt Deatherage

March 1991

Written by: Keith Rollin

November 1988

This Technical Note corrects an error in the record used for GetPenState and SetPenState.

Changes since November 1988: This note is now obsolete.

This Note formerly described an error in the pen state record in Volume 1 of the Apple IIgs Toolbox Reference. This error is corrected on page 43-2 of Volume 3 of the Toolbox Reference.

Further Reference

- o Apple IIgs Toolbox Reference, Volume 3

Apple IIGS
#45: Parameters for GetFrameColor

Revised by: Matt Deatherage
Written by: Dan Oliver

September 1989
November 1988

This Technical Note formerly attempted to correct the description of the parameters passed to and returned from the routine GetFrameColor in the Window Manager chapter of the Apple IIGS Toolbox Reference. This call works as documented since System Software 3.2; therefore, former versions of this Note were incorrect.

Changes since November 1988: Corrected our error. Sorry for any inconvenience.

This Note formerly stated the following: "The Apple IIGS Toolbox Reference, Volume 2 incorrectly describes the parameters passed to and returned from GetFrameColor on page 25-57."

However, this is incorrect. Beginning with System Software 3.2, GetFrameColor works as documented in the Apple IIGS Toolbox Reference, Volume 2. Prior to System Software 3.2, the call did not work at all. We apologize for any inconvenience this confusion may have caused.

Further Reference

- o Apple IIGS Toolbox Reference, Volume 2

Apple IIGS
#46: DrawPicture Data Format

Written by: Jeff Erickson & Keith Rollin

November 1988

This Technical Note describes the internal format of the QuickDraw II picture data structure.

This Technical Note presents the internal format of the QuickDraw II picture data structure for informational purposes only. You should not use this information to write your own bottleneck procedures; the only routines which should create and read PICT format files are those provided in QuickDraw II. If we added new objects to the picture definition, your program would not operate on new pictures. This Note documents this information for debugging purposes only.

Picture Data Structure Definition

Pictures are stored in memory in the following format:

They begin with a WORD which indicates the mode of the port which was used to record when the picture was created. This information is useful when the picture is played back, possibly in a different graphics mode.

Following the WORD is a RECT which indicates the frame of the picture and is used for scaling when you redraw the picture. Following the RECT is the version number of this PICT format, then a series of word-sized opcodes which describe the sequences of QuickDraw II commands that were used to create the picture.

Name	Description	Size (bytes)
pictSCB	picture's scan line control byte	2 (high byte = 0)
picFrame	picture's boundary rectangle	8
version	picture version	2 (Currently \$8211)
opcode	operation code	2
<data>	operation data	variable, depending on opcode
:		
opcode	operation code	2
<data>	operation data	variable, depending on opcode

Opcodes

As mentioned above, pictures are described by a series of opcodes which are used to record the QuickDraw II commands that created the picture. These opcodes are two bytes long and are usually followed by a number of parameters.

All currently defined opcodes and their parameters are listed below. Any opcodes not listed here are reserved.

Opcode	Name	Description	Parm Bytes	Parameter Description
\$0000	NOP	no operation	0	none
\$0001	ClipRgn	clip to a region	[region size]	region
\$0002	BkPat	background pattern	32	background pattern (8x8)

\$0003	TxFont	text font	4	pixels) Font Manager font ID (long)
\$0004	TxFace	text face	2	text face (word)
\$0005	TxMode	text mode	2	text mode (word)
\$0006	SpExtra	space extra	4	space extra (fixed)
\$0007	PnSize	pen size	4	pen size (point)
\$0008	PnMode	pen mode	2	pen mode (word)
\$0009	PnPat	pen pattern	32	pen pattern (8x8 pixels)
\$000A	FillPat	fill pattern	32	fill pattern (8x8 pixels)
\$000B	OvSize	oval size	4	oval size (point)
\$000C	Origin	origin	4	origin (point)
\$000D	TxSize	text size	2	text size (word)
\$000E	FGColor	foreground color	2	color (word)
\$000F	BGColor	background color	2	color (word)
\$XX11	Version	version	0	none: high byte=version (currently \$82)
\$0012	ChExtra	character extra	4	char. extra (fixed)
\$0013	PnMask	pen mask	8	mask (8 bytes)
\$0014	ArcRot	arc rot	2	Reserved (related to things drawn with patterns). (word)
\$0015	FontFlags	font flags	2	font flags (word)
\$0020	Line	line	8	pnLoc (point), newPt (point)
\$0021	LineFrom	line from pen loc.	4	newPt (point)
\$0022	ShortLine	short line	6	pnLoc (point), dv, dh (signed bytes)
\$0023	ShortLFrom	ditto from pen loc	2	dv, dh (signed bytes)
\$0028	LongText	long text	5+text	txLoc (point), count (byte), text
\$0029	DHText	hor. offset text	2+text	dh (unsigned byte), count (byte), text
\$002A	DVText	vert. offset text	2+text	dv (unsigned byte), count (byte), text
\$002B	DHDVText	offset text	3+text	dv, dh (unsigned bytes), count (byte), text
\$002C	RealLongText	very long text	6+text	txLoc (point), count (word), text

Opcodes between \$0030 and \$008C are a combination of a graphic verb and a graphic object, as listed below (where "V" stands for the graphic verb, and "X" is a stands for the graphic object). For example, \$0069 means PaintSameArc, and is followed by two one-word parameters.

Graphic Verbs:

\$00X0	Frame...	frame something	[Specific to object type see below.]
\$00X1	Paint...	paint something	
\$00X2	Erase...	erase something	
\$00X3	Invert...	invert something	
\$00X4	Fill...	fill something	
\$00XV+8	...Same...	draw same thing somehow	[See below; {braced} parms do not appear.]

Graphic Objects:

\$003V	...Rect	draw a rectangle somehow	8 (0 if - SameRect) {rect (2 points)}
\$004V	...RRect	draw a round rect somehow	8 (0) {rect (2 points)}

\$005V	...Oval	draw an oval somehow	8 (0) {rect (2 points)}
\$006V	...Arc	draw an arc somehow	12 (4) {rect (2 points)}, start, arc angle (words)
\$007V	...Poly	draw a polygon somehow	[polygon size] (0) {polygon}
\$008V	...Rgn	draw a region somehow	[region size] (0) {region}
\$0090	BitsRect	copybits, rect clipped	variable* (see below, but without maskRgn)
\$0091	BitsRgn	copybits, rgn clipped	variable* (see below)
\$00A1	LongComment	long comment	4+data kind (word), size (word), data

*Bits... data:

origSCB	original scan line control byte	2	SCB (word -- high byte = 0)
BWvsColor	black and white vs. color	2	reserved (word)
width	width of pixel image in bytes	2	width (word)
boundsRect	bounds rectangle	8	rect (2 points)
srcRect	source rectangle	8	rect (2 points)
destRect	destination rectangle	8	rect (2 points)
mode	transfer mode	2	pen mode (word)
maskRgn	mask region (BitsRgn ONLY!)	[region size]	region
pixData	pixel image	[pixdata size]	width* (bounds.bottom- bounds.top)

Differences Between IIGS Pictures and Macintosh Pictures

1. QuickDraw II pictures are modeled after PICT2 on the Macintosh, which use two bytes for its opcodes and data (the exception to this is the \$11 (version) opcode, which is followed by a one-byte parameter). Macintosh PICT 1.0 formats, which use one-byte opcodes, would have to undergo extensive modifications to be displayed on the IIGS.
2. There is no EndOfPicture opcode on the IIGS as there is on the Macintosh. Also, the first word of the picture is a pictSCB, not the length of the picture. The picture size is determined solely by the size of the handle on the IIGS. There is also no picture header on the IIGS as on the Macintosh.
3. The number sex of the Macintosh is opposite that of the Apple IIGS. The Macintosh stores the high bytes of words and long words first, whereas the IIGS stores the low byte first.
4. The following Macintosh picture opcodes are not available on the IIGS: txRatio, PackBitsRect, PackBitsRgn, shortComment, EndOfPicture.
5. QuickDraw II defines the following opcodes that the Macintosh does not: ChExtra (\$12), PnMask (\$13), ArcRot (\$14), FontFlags (\$15), and RealLongText (\$2C).

Notes on the Interpretation of IIGS Pictures

- o The state of the pen, the clip region, various patterns and colors, and the origin of the current port is saved before a picture is drawn, and restored afterwards. The current port is set up in a default state equivalent to that of a newly created port just before drawing begins. Picture opcodes act just like their QuickDraw II tool counterparts, with a few exceptions.
- o Two pen locations are tracked as the picture is drawn, one for lines and one for text. Thus, LineFrom always draws from the end of the last line, regardless of any intermediate text opcodes.

- o Text calls do not change the position of the "text pen," as do normal QuickDraw II text calls. Thus, if a picture contains two lines of text, the second one directly below the first, the second will be stored using a DVtext opcode.
- o DrawPicture performs considerable setup before it draws pictures. Among other things, it calls InstallFont, which is a Font Manager call. If you are going to support pictures in your application, you should load and start the Font Manager.

Further Reference

- o Apple IIGS Toolbox Reference, Volume 2

Apple IIGS

#47: What SetDataSize Does

Written by: Keith Rollin

November 1988

This Technical Note clears up any ambiguity in the description of the SetDataSize call.

The Apple IIGS supports windows that contain scroll bars in their frames. These scroll bars are handled by TaskMaster and differ from Macintosh scroll bars in that the size of the "thumb" or "elevator" is used to indicate the size of the visible area of the document in relation to the total size of the document (the "data size"). Initially, the visible size and the data size are defined by the parameter list passed to NewWindow; however, either of these can be changed by SizeWindow and SetDataSize, respectively.

SetDataSize is used to not only change the range of scrolling allowed, but also to redraw the size of the thumb to reflect the fact that the data size has changed with respect to the visible area. However, page 25-97 of the Apple IIGS Toolbox Reference contains the following description of SetDataSize:

"Sets the height and width of the data area of a specified window. Setting these values will not change the scroll bars or generate update events."

When the manual states that SetDataSize "will not change the scroll bars," it is referring to the location, or value, of the thumb. Assume a situation where you have a word processor that scrolls the page using TaskMaster scroll bars. If you delete a range of text, you would also shorten the entire size of the document. Calling SetDataSize to reflect that would indeed change the size of the thumb, but it would not change its location. If you were already scrolled to the bottom of the document when you called SetDataSize, the thumb would become larger (to reflect the fact the the total data size became smaller with respect to the visible data size) and overwrite the down arrow of the scroll bar. To prevent this situation from occurring, you should also change the origin of the window with SetContentOrigin before calling SetDataSize.

Further Reference

- o Apple IIGS Toolbox Reference, Volume 2

Apple II GS
#48: All About AlertWindow

Revised by: Dave Lyons
Written by: Dan Oliver & Keith Rollin

July 1989
November 1988

This Technical Note documents a new call in the Window Manager which eases the creation of Alert windows.

Changes since July 1989: The information on AlertWindow formerly found in this Note has been updated and is now included in the Apple IIgs Toolbox Reference, Volume 3.

The information on AlertWindow formerly found in this Note has been updated and is now included in the Apple IIgs Toolbox Reference, Volume 3. This Window Manager call was first introduced in System Software 3.2.

Apple IIGS
#49: Rebooting (Really)

Revised by: Matt Deatherage
Written by: Matt Deatherage & Jim Merritt

January 1989
November 1988

This Technical Note discusses rebooting the Apple IIGS from software.
Changed since November 1988: Corrected two assembly-language instructions in the FROMNATV routine in the example code.

In days gone by, many Apple II applications had a Quit menu option. Unfortunately, a large number of these simply rebooted the machine. Today, this is far from desirable. Even with the advantages of GS/OS-reduced booting time (around 34 seconds with an Apple 3.5 Drive), waiting for the operating system to reload, as well as wiping out any ongoing tasks by desk accessories (such as an alarm clock) makes the standard ProDOS 8 or GS/OS QUIT call much more attractive.

However, there are still instances where an application may wish to require the user to reboot. A common example might be a game. The game might use GS/OS in a completely standard way, but if you QUIT from the program GS/OS booted into, you will be returned to the same program. Since most applications will boot into the Finder, this is not a widespread problem. However, the Finder must also provide the reboot option, and alternate program selector applications may wish to provide this functionality as well.

The Easy Way

GS/OS provides a mechanism for rebooting with the OSShutdown call. This call, documented in GS/OS Reference, Volume 1, will either reboot the system (after first shutting down all loaded and generated drivers and closing all open sessions) or will shut down everything and present a dialog box which states, "You may now power down your Apple IIGS safely." A Restart button is provided which allows the user to reboot without pressing Control-Open Apple-Reset .

Note: When using System Disk 4.0, if the Window Manager is active when you issue the OSShutdown call, there must be at least one open window; it need not be visible, but it must be open. This will be fixed in the next revision of GS/OS.

The OSShutdown call also provides a way to resize the internal RAM disk (named /RAM5 by default). Most programs have absolutely no need to use this mechanism, and should avoid it whenever possible. A notable exception would be a third-party RAM disk utility which uses a battery backup, which may need to make changes which require resizing the RAM disk. Of course, such a utility should ask the user to ensure that erasing the RAM disk content is acceptable. Resizing the RAM disk is only possible when using the OSShutdown call; any other method you may be using to accomplish this function from software will break in the future.

If you are using GS/OS, you should always use OSShutdown. You must not reboot the system in any other fashion. The OSShutdown mechanism provides a convenient and supported way to restart or shut down the system. Doing it another way can easily cause a loss of data.

The Hard Way

Programs not using GS/OS have a little more work to do. The supported non-GS/OS method of rebooting is similar to the method used on 8-bit machines: change the value of POWERUP (\$00/03F4) and do a long jump to RESET (\$FA62). However, there are a few catches:

1. The jump must be made in emulation mode.
2. Interrupts must be disabled.
3. The data bank register must be set to zero.
4. The direct page must be zero.
5. ROM firmware must be visible in the memory map.
6. Internal interrupt sources (such as the ones for AppleTalk) must be shut down.

Simply disabling interrupts without shutting down AppleTalk interrupt sources inside the system will cause the system to hang when the jump to RESET is made. Turning off these internal interrupt sources is accomplished by changing softswitch values at \$C039 (SCCAREG), \$C041 (INTEN), and \$C047 (CLRVBLINT).

The following code example demonstrates the correct method:

```
POWRUP      equ    $0003F4    ;the power-up byte in bank zero
STATEREG    equ    $C068      ;ROM/RAM state register
CLRVBLINT   equ    $C047      ;clear VBL interrupt flags register
INTEN       equ    $C041      ;interrupt enable register
SCCAREG     equ    $C039      ;SCC register
RESET       equ    $00FA62    ;ROM reset entry point
;
FROMNATV    anop           ;enter here from native mode
            sei            ;disable interrupts
            pea    0
            pea    0        ;push four zero bytes on the stack
            plb            ;pull data bank register
            plb            ;(twice to balance the stack)
            pld            ;pull 16-bit data bank register
            sec
            xce            ;go into emulation mode
            longa    off
            longi    off
FROMEMUL    anop           ;enter here from emulation mode
            sei            ;disable interrupts for people entering here
            dec    POWRUP    ;invalidate the power up byte
            lda    #$0C      ;ROM parameters
            sta    STATEREG  ;swap in the ROM and everything else out
            stz    CLRVBLINT ;clear VBL interrupts
            stz    INTEN     ;turn off internal interrupt sources
            lda    #$09
            sta    SCCAREG   ;shut down SCC interrupt sources
            lda    #$C0
            sta    SCCAREG
            jml    RESET     ;and off we go into the wild blue yonder
```

These methods of restarting the system are presented for those applications that absolutely must do so. Rebooting is not a suggested way of ending an application and the techniques described in this Note should be used with extreme caution.

Further Reference

- o Apple IIGS Firmware Reference
- o GS/OS Reference, Volume 1

Apple IIGS

#50: Extended Serial Interface Error Handling

Written by: Dan Strnad

January 1989

This Technical Note discusses error reporting by the Extended Serial Interface.

For Apple IIGS ROM 01, the Extended Serial Interface does not return the error condition in the carry bit. Programs using the Extended Serial Interface should check for a non-zero result value in the result code rather than the carry bit to determine if an error has occurred. The following eight-bit APW code demonstrates this error checking using the SetDTR command. The SetDTR routine zeros the result bytes if no error has occurred.

```
                LONGA    OFF                ;PREPARE ASSEMBLER FOR EMULATION MODE
                LONGI    OFF
                65C02    ON
                KEEP     SETDTR2
                START
SLOT            EQU      $01
                SEC      ;SET EMULATION MODE
                XCE
                JMP      BEGIN
CMDLST          DC       H'03'             ;PARAMETER COUNT
                DC       H'0B'             ;SETDTR COMMAND CODE
RESLT           DC       I'0'             ;RESULT CODE (OUTPUT)
DTRSTAT         DC       I'0'             ;BIT 7 IS STATE OF DTR (INPUT)
BEGIN           LDA      #SLOT             ;COMPUTE $CN VALUE TO BE USED
                ORA      #$C0
                STA      OFFSET+2         ;MODIFY INSTRUCTIONS LOADING OFFSETS
                STA      XOFFSET+2
                STA      ICALL+2          ;MODIFY INSTRUCTIONS CALLING FIRMWARE
                STA      XCALL+2
IOFFSET         LDA      $C00D             ;THIS INSTRUCTION MODIFIED AT RUNTIME
                STA      ICALL+1         ;MODIFY JSR TO INIT
XOFFSET         LDA      $C012             ;THIS INSTRUCTION MODIFIED AT RUNTIME
                STA      XCALL+1         ;MODIFY JSR TO EXTENDED SERIAL INTERFACE
ICALL           JSR      $C000             ;THIS INSTRUCTION MODIFIED AT RUNTIME
                LDA      #<CMDLST        ;LOW BYTE OF COMMAND LIST
                LDX      #>CMDLST        ;HIGH BYTE OF COMMAND LIST
                LDY      #0               ;24-BIT ADDRESS NOT USED BY 8-BIT PROGRAM
XCALL           JSR      $C000             ;THIS INSTRUCTION MODIFIED AT RUNTIME
                LDA      RESLT           ;DID AN ERROR OCCUR?
                BNE      ERROR           ;YES- HANDLE THE ERROR
                ...
ERROR           ...
                END
```

Apple IIgs

#51: How to Avoid Running Out of Memory

Revised by: Dave Lyons
Written by: Eric Soldan

May 1992
January 1989

This Technical Note discusses handling nearly-out-of-memory situations when working with the IIgs tools.

CHANGES SINCE SEPTEMBER 1990: Added discussion of an Out-of-memory routine problem fixed in System 6.0.

INTRODUCTION

Running out of memory is a concern for most every application. Working with the Toolbox makes monitoring this situation a little more difficult since your application is not the only one allocating memory.

Low-level toolbox functions (for example, QuickDraw II calls) require that a 16K block of memory be allocatable, while high-level routines (for example, the Window Manager) require that a 32K block of memory be allocatable. Apple does not guarantee that toolbox functions behave reasonably if there is less memory available, and the tools are not stress-tested with less than the minimum required memory available.

Since the toolbox assumes reasonable memory-allocation requests succeed, just waiting for an out-of-memory error is not adequate memory management. To make your application work reliably in low-memory situations, you need a method of ensuring that the toolbox gets memory when it needs it. This Note describes two approaches.

HOW MUCH MEMORY CAN BE ALLOCATED

There's no way to tell how much memory can be allocated without actually trying to allocate it.

MaxBlock tells you the size of the largest single free block, but this doesn't take into account purgeable blocks, compaction, and out-of-memory routines (see Apple IIgs Toolbox Reference, volume 3). FreeMem and RealFreeMem cannot tell you how badly fragmented the memory is, and they do not take into account out-of-memory routines.

A SUGGESTED METHOD

A method of checking for a nearly-out-of-memory condition is to have your own purgeable handle just for this task. If the handle has not been purged, then you have plenty of memory for the toolbox, and in the worst case, the toolbox purges your handle if it needs the RAM.

The less often your purgeable handle gets purged, the better performance you get in nearly-out-of-memory situations. Therefore, you should arrange for other purgeable memory, not necessarily belonging to your application, to be purged before your handle. For example, you want dormant applications to be purged, rather than having your handle get repeatedly purged and reallocated.

So the purge level of this handle should be one.

The check to see if a handle has been purged is very fast. If it has been purged, you have to try to reallocate it. Reallocating a handle is not a fast process, so the fewer times the handle is purged, the faster the check is and the better your performance. Unless you are in a nearly-out-of-memory situation, the handle should not be purged at all, and you should have virtually no overhead for this process.

This technique can be implemented as follows:

```
appStart
;
; Somewhere at start, create a purgeable handle of size N,
; called "loMemHndl", purge level 1.
;
                rts

*****
;
; Here's an example of checking for nearly-out-of-memory:
;
                jsr    preCheckLoMem
                bcc    goForIt
                bcs    HandleError          ;Handle errors appropriately.
goForIt         (_ToolboxCall[s])          ;Make as many as needed.
;
; Here you can make your toolbox calls. Since you prechecked
; for nearly-out-of-memory conditions, you should have no memory
; errors at this point.
;
; You could also check after calls, as shown here:
;
                (_ToolboxCall)
                jsr    checkLoMem          ;Call this to see if low.
                bcc    noError
                bcs    HandleError          ;Take care of errors.

noError         jsr    lifeIsGood
                .
                .
                .
                rts

*****
;
; Here are some sample routines to check for the nearly-out-of-
; memory condition.
;
checkLoMem      bcs    retErr
preCheckLoMem   lda    [loMemHndl]
                ldy    #2
                ora    [loMemHndl],y
                beq    gotPurged
                lda    #0
                clc
                rts

gotPurged       (Try reallocating it into loMemHndl, purge level 1.)
                (If you can't, you will get a $0201 error. You may wish to
                return the $201 error, or you may wish to change it into
                your own error code.)
```

```
i
retErr          rts          ;This is a single exit point
                                ;whether errors were present
                                ;or not.
```

You can determine the size of this purgeable handle, but like determining what size stack is adequate for an application, there is no single "right" answer. There are different considerations for size of the purgeable handle for each application, and these may change during the development process. Use your best judgement, keeping in mind that high-level toolbox routines require a 32K block.

AN ALTERNATIVE

For better control over when your handle is purged or disposed, you can write an out-of-memory routine as described in the Memory Manager chapter of Apple IIgs Toolbox Reference, volume 3. Out-of-memory routines have the opportunity to free up memory before or after the Memory Manager attempts to purge purgeable handles, and this manual contains a sample of such a routine.

NOTE : If your Out-of-memory routine frees up memory on the second pass, there is a problem with the Memory Manager in System Software 5.0 through 5.0.4 that may affect you. If your routine frees enough bytes on the second pass, but the Memory Manager still cannot complete the request it is working on, it can hang for a couple of minutes and then crash. This is fixed in System 6.0.

Further Reference:

- o Apple IIgs Toolbox Reference, Volumes 1-3

Apple IIgs
#52: Loading and Special Memory

Revised by: Dave Lyons
Written by: Eric Soldan

May 1992
January 1989

This Technical Note discusses strategies for preventing applications from loading into special memory.

CHANGES SINCE JULY 1989: Noted that the System 6 Loader always tries non-special memory before special memory.

The System 6.0 Loader always tries to load segments into non-special memory before allowing them to load into special memory. The rest of this Note is useful if your application does not require System 6, or if you need an example of an initialization segment.

The System Loader loads your application starting at the lowest memory location possible. If you allow your program to load into special memory, the Loader first tries bank \$01. If your program cannot load into special memory, it starts at bank \$02. Either way, the Loader progresses to higher banks, and eventually, it may even try loading into bank \$E1, which contains the super hi-res screen.

The problem with allowing your application to load into special memory is that the super hi-res screen is part of special memory. If you have a desktop application, part of your application may load into the super hi-res screen, and when you try to start QuickDraw II, it fails because the screen memory is already allocated.

When QuickDraw II fails because your program loaded into the SHR screen, it seems reasonable to assume that the Loader put your program there because it needed the RAM which special memory provides. This logic seems to make sense, but it is not completely reliable. The Loader (in System Software earlier than 6.0) tries to put your program into special memory before it tries purging dormant applications. This means that the more programs that run from the Finder that set the GS/OS or ProDOS 16 "restartable from memory" bit, the more likely it is that the next application launched that can load into special memory will load into the super hi-res screen.

For this reason, it is important not to let your application load into special memory, or at least not load into the super hi-res screen. If your application is not allowed to load into special memory, then the Loader will purge other dormant applications to make space for yours. One way to accomplish this is when linking your application. You can set the "no special memory" bit in the OMF KIND field of applications using OMF 2.0 or later, but this also prohibits your application from using bank \$01.

Another way to avoid loading into the super hi-res screen is to have your initial segment allocate the super hi-res screen. You can accomplish this by starting QuickDraw II in your initial segment, then the rest of your program cannot load into the already-allocated super hi-res screen. This strategy could fail if the initial segment loaded into the super hi-res screen, but this is very unlikely and can be prevented by flagging the initial segment to only load into non-special memory. You can do this by setting the "no special memory" bit in the KIND field only for the initial segment.

Here's an example of such an initial segment in MPW IIgs format:

```
*****
*
* You may wish to do this stuff in the initial segment of your
* application.  The initial segment should be set so that it does not
* load into special memory, or else it is possible that it would load
* into the super hi-res screen.  If this occurred, then QuickDraw II would
* not be able to be started.
*
* Once QuickDraw II is started, the super hi-res screen is taken,
* therefore the rest of the application can not load into it.  Therefore,
* special memory is generally an acceptable place for the rest of the
* application to load, since the special memory needed for the screen
* is already taken.
*
* If the performance of your application would be adversely affected
* by memory fragmentation, then you should also consider purging
* other dormant applications and dormant tools, and then compacting
* memory.  This will prevent fragmentation as much as possible
* while your application is loading.  It also has the cost of longer
* startup time since some tools may have to be reloaded.  This is the
* only way to be sure that tools that you don't want are removed
* from memory before the rest of your application tries to load
* around them.
*
* The Finder is a dormant application when your application is
* launched.  This will cause the Finder to be thrown out of memory,
* and it will have to be reloaded when your application is quit.
*
*****
```

case on

```
include 'e16.memory'
include 'm16.memory'
include 'm16.quickdraw'
```

```
screenMode      equ    $80
AppMaxWidth     equ    160                ;Double this and your application
                                           ;will print in BetterText mode.
```

initialScreen PROC

```
myID             equ    1                ;long
zpagehndl       equ    myID+4           ;long
```

```
stkAfterLocals equ    zpagehndl+4
```

```
directReg       equ    stkAfterLocals
retAddr         equ    directReg+2
passedParms     equ    retAddr+3
```

```
    phd                ;Set up stack frame.
    tsc
    sec
    sbc    #stkAfterLocals-1
    tcs
    tcd
    pha
```

```

_MMStartUp
pla
sta myID ;Get the userI
pha
_HLockAll ;Lock down the rest of ourselves, in
;case we are being restarted. The
;loader does not prelock down stuff,
;so we would be disposing of the rest
;of ourselves.

pea $1000
_PurgeAll ;Purge other dormant applications.
;This is optional.

pea $4000
_PurgeAll ;Purge dormant tools.
;This is optional.

_CompactMem ;Clean up memory. This is advised.

pha ;Make direct space for QuickDraw.
pha
pea $300>>16 ;Hi-byte of $300 address.
pea $300
pei myID
pea attrLocked+attrFixed+attrPage+attrBank
lda #0
pha
pha
_NewHandle
plx
stx zpagehndl
plx
stx zpagehndl+2
bcc @a
ERRORDEATH 'Out of bank 0 memory'

@a
lda zpagehndl
sta >qdstarthndl ;Used for disposing handle at shutdown.
txa
sta >qdstarthndl+2
lda [zpagehndl] ;Start up QuickDraw. This protects
pha ;screen RAM from the rest of the
pea screenMode ;application loading into it.
pea AppMaxWidth
pei myID
_QDStartUp
bcc @b
ERRORDEATH 'Can''t start up QuickDraw'

@b
;Do title screen here.

tsc
clc
adc #stkAfterLocals-1
tcs
pld
rtl

qdstarthndl dc.l 0

ENDP
END

```

Further Reference:

- o GS/OS Reference, Volume I
- o MPW IIgs Tools Reference
- o APW Assembler Reference

Apple IIgs

#53: Desk Accessories and Tools

Revised by: Dave "Out of Phase" Lyons

May 1992

Written by: Matt Deatherage, Jim Mensch, & Dave Lyons

March 1989

This Technical Note discusses compatibility issues that can arise between desk accessories and applications. Where possible, it presents solutions.

CHANGES SINCE MARCH 1991: Updated information about QuickDraw Auxiliary and StartUpTools for System 6.0.

This Note presents guidelines to help applications and desk accessories work together smoothly.

TOOL SETS

The greatest conflict between applications and desk accessories, especially NDAs, is the use of system tool sets. The Apple IIgs Toolbox Reference, Volume 1, defines the minimum collection of tools sets available to an NDA. The Desk Manager requires that an application start the following tool sets before calling DeskStartUp:

- Tool Locator (#1)
- Memory Manager (#2)
- Miscellaneous Tools (#3)
- QuickDraw II (#4)
- Event Manager (#6)
- Window Manager (#14)
- Menu Manager (#15)
- Control Manager (#16)
- LineEdit (#20)
- Dialog Manager (#21)
- Scrap Manager (#22)

NDAs may assume that these tools are all present and running, so they do not need to check for their presence. NDAs can also use the following tool sets without special consideration for starting them up: Desk Manager, Scheduler, Apple Desktop Bus, and Integer Math.

In addition to the tool sets applications must start to support NDAs, Apple recommends that applications start the following tools:

- QuickDraw Auxiliary (#18) (see discussion under QuickDraw Auxiliary)
- Font Manager (#27)

These two additional tools are so widely used by desk accessories that they should be present. NDAs may not assume their presence, but it is reasonable to write an NDA that checks for them, with the assumption that they usually turn out to be available.

NDA GUIDELINES

WHICH TOOL SETS CAN AN NDA USE?

- o In general, NDAs can use the tool sets which have already been started up by the host application, even tools that are not guaranteed to be started up. Using other tool sets is trickier (see below).
- o In general, NDAs should not start up tools that are already started up. (The Resource Manager is an exception.)
- o The Resource Manager must be started separately by each client. See Apple IIgs Technical Note #71 for detailed information on using the Resource Manager from an NDA.
- o Sound tools are an exception to the rule of freely using a tool which is already started. See the section "Sound Tools" sections later in this note.
- o Some tool sets are easily started up each time they are needed, if they are not already present.

Standard File is an excellent example. If an NDA needs to use Standard File, it should check to see if the tool is already running. If it is not running, the NDA must use LoadOneTool to load it, then it must allocate a page of direct-page space and start the tool before using it. When finished with the tool, the NDA must shut it down, deallocate the direct-page space, and unload it with UnloadOneTool. (A tool is already running if its xxxStatus function returns TRUE and does not return an error.)

The important thing here is that the NDA shuts down Standard File immediately after using it, if it was not already started. This does not cause conflicts with the host application or with other NDAs.

Note that by pre-initializing the result space of an xxxStatus call to zero, you can avoid caring whether the tool is present but not started or simply not present.

```

    pea $0000
    _SFStatus
    pla                                ;A is nonzero if Standard File is started

```

From a high-level language, you may not be able to pre-initialize the result space. Instead, you need something like the C statement:

```
StdFileActive = ( SFStatus() && !_toolErr);
```

or the Pascal statement:

```
StdFileActive := (SFStatus<>0) AND (ToolErrorNum=0);
```

- o It is impractical or impossible to start up certain tool sets each time they are needed. These include the Font Manager, Scrap Manager, and Text Edit.

If an NDA needs to start up a tool and keep it started while letting the application continue to run, things get interesting. (There is a risk that the host application will later try to start up the tool set itself and not be able to deal with the tool already being started.)

In practice, the safest thing you can do for a tool you need to leave running is:

--When your NDA is opened, check the tool set's status. If it is not available, use LoadOneTool, allocate any needed direct-page space, start up the tool set, and set a flag indicating that your NDA started the tool set.

--When your NDA's Init routine is called at DeskShutDown time (Accumulator equal to zero), check the flag set above. If your NDA started a tool set, shut it down, dispose of any direct-page space you allocated for it, and call UnloadOneTool.

(Keep in mind that your NDA can be opened and closed many times before DeskShutDown is called when the application finally quits. If you have started a tool and set a flag on an open, be sure not to disturb the flag on a future open, when the tool is already available because you started it! You still need to shut it down at DeskShutDown time.)

--Do not shut down tool sets when your NDA is closed. To see why, consider what would happen if two NDAs just like yours were used at the same time. If the NDAs were closed in any other than the exact opposite order they were opened, some NDAs would have tool sets shut down from underneath them.

STARTUPTOOLS

- o StartUpTools in System Software 5.0.4 and earlier is designed to be called only by an application, not a desk accessory. Unexpected things happen if your NDA calls StartUpTools (for example, you may get a second copy of the application's resource fork open in your NDA's private resource search path; this wastes RAM and can interfere with an application's attempt to write to its own resource fork).
- o See the System 6.0 Toolbox documentation for information on using StartUpTools from an NDA. There are new flag bits you need to know about.

TLSTARTUP AND TLSHUTDOWN

- o Do not call TLStartUp or TLShutDown from a desk accessory.
- o You may call MMStartUp at any time to get your desk accessory's own memory ID. This does not allocate a new ID; it just tells you what ID you already have (it returns the memory ID of the block the MMStartUp call is made from).

USER TOOL SETS BELONG TO THE APPLICATION

- o A desk accessory must not install user tool sets, because there is no arbitration of user tool set numbers. User tool sets are the sole property of the current application.

A desk accessory should not call user tool sets even if it determines that the host application has installed a certain tool set, because that limits future system software options. For example, consider a hypothetical multiple-application environment. If DAs call user tool sets and the system automatically switches between separate collections of user tool sets, there would be no way for the system to know which set to switch in before giving control to a desk accessory.

BANK ZERO MEMORY AND ERROR \$0201

- o If you get error \$0201 (unable to allocate memory block) while trying to launch a ProDOS 8 application, it is probably because your NDA

allocated some memory in bank 0 or bank 1 and failed to dispose of it at DeskShutDown time (when the NDA's Init routine is called with the accumulator equal to zero). GS/OS needs to allocate all of this memory for ProDOS 8 to use.

QUICKDRAW AUXILIARY

- o In System 6.0 and later, QuickDraw Auxiliary is always available to an NDA, because the Window Manager automatically loads and starts QuickDraw Auxiliary (because it's needed for AlertWindow, for example). To prevent problems, duplicate QDAuxStartUp and QDAuxShutDown calls are tolerated, and QDShutDown automatically calls QDAuxShutDown.
- o Before System 6.0, starting QuickDraw Auxiliary when the application has not started it can be a problem. An application that correctly implements switching between 320 and 640 mode calls QDShutDown and QDStartUp. QuickDraw Auxiliary depends heavily on QuickDraw, and restarting QuickDraw while QuickDraw Auxiliary is active will fry big-time.

SOUND TOOLS

- o A desk accessory cannot use any of the sound tools if they are already started. This is contrary to the rule for other tool sets, but it is required because there is no memory management of the sound RAM (or "DOC RAM"). If the Sound Tools (#8) are started, the application has exclusive control of the 64K DOC RAM used to play sounds. Anything your desk accessory might put there could overwrite information the application needs.

Saving and restoring DOC RAM around desk accessory usage is not sufficient. Many of the sound functions are interrupt driven, altering the contents of DOC RAM only during sound interrupts, so your desk accessory might attempt to replace parts of DOC RAM which are being played. Since there is no memory management of DOC RAM, desk accessories must avoid the sound functions of the IIgs if the application is already using them.

APPLICATION GUIDELINES

For best compatibility with NDAs, applications should follow the following guidelines.

- o Be careful about when your application starts and shuts down tools. A highly compatible approach is to start tools at the beginning of your application and leave them started. For certain tools, like Standard File, it is reasonable to load and start the tool set each time it's needed (you may want to check whether it's already started, in case some impolite NDA started Standard File and left it started).

Note that UnloadOneTool followed later by LoadOneTool does not necessarily cause disk access or ask the user to insert the boot disk. UnloadOneTool calls UserShutDown to put the tool set into "zombie" state, where it can be restarted from memory if none of its segments have been purged. Unloading tools while they aren't in use is a Good Thing--if the user has plenty of RAM, there's no noticeable performance hit, but if RAM space is tight then doing extra disk access still is preferable to actually running out of memory.

For maximum compatibility, an application should not shut down any tools which were ever active when it called SystemTask or TaskMaster (until quitting time, of course, when it shuts down everything, starting with the Desk Manager). The application can start more tools, but it should not shut down those which are already active.

If your application is going to start a tool and not keep it started, use it and then shut it down immediately, without allowing desk accessories to be opened during that time.

- o Don't just start the Scrap Manager--use it! Many desk accessories support cutting and pasting to exchange text and pictures with your application, but they can do it only if you use the Scrap Manager. If you have a need for your own private scrap internally, you should still also use the Scrap Manager to exchange text and pictures with other applications and DAs.
- o Provide an Edit menu, and when an NDA window comes to the front enable the menu and the Undo, Cut, Copy, Paste, and Clear items.
- o Applications should never make a Close call with reference number zero at file level zero. (If you need to use Close with reference number zero, use GetLevel and SetLevel to avoid closing files you did not open.)

DAs written recently can open their files at an internal file level, as documented in GS/OS Technical Note #13, but applications still need to avoid closing all files at level zero for compatibility with older desk accessories.

- o An application with some memory to spare can save NDAs time by providing them the additional tools which they are most likely to use.

The most common tools which desk accessories require besides those available in the standard Desk Manager set are QuickDraw Auxiliary (#18), the Print Manager (#19), Standard File (#23), the Font Manager (#27), and the List Manager (#28).

- o When you call TaskMaster or GetNextEvent, or EventAvail, be sure bit 10 is turned on in the event mask, to enable "desk accessory" events. If you turn this bit off, users will not be able to get to the Classic Desk Accessory menu by pressing Apple-Ctrl-ESC.

CDA GUIDELINES

- o CDAs are nearly always modal, but by using the HeartBeat interrupt queue or other mechanisms, they can get control when the user is no longer "in" the CDA. The list of guaranteed tools for NDAs does not apply to CDAs, and CDAs must be prepared to deal with the ProDOS 8 environment as well as GS/OS.
- o Under ProDOS 8, a CDA will not be able to allocate any bank 0 space through the Memory Manager; it can only use page 0 and page 1 safely (the stack is in page 1).
- o Do not call TLStartUp or TLShutDown from a desk accessory.
- o You may call MMStartUp at any time to get your DA's own memory ID. This does not allocate a new ID; it just tells you what ID you already have (it returns the memory ID of the block the MMStartUp call is made from).

Further Reference:

- o Apple IIgs Toolbox Reference
- o Programmer's Introduction to the Apple IIgs
- o Apple IIgs Technical Note #71, Desk Accessory Tips and Techniques
- o Apple IIgs Technical Note #83, Resource Manager Stuff

Apple IIgs
#54: MIDI Drivers

Revised by: Matt Deatherage
Written by: Jim Mensch

November 1990
May 1989

This Technical Note describes how to write a driver for use with the Apple IIgs MIDI tools.

Changes since May 1989: Noted that MIDI drivers also work with the MIDI Synth tool.

Apple ships two drivers with the MIDI tool set, APPLE.MIDI and CARD6850.MIDI, respectively. These drivers are adequate for almost all MIDI hardware currently on the market for the Apple IIgs; however, if your hardware is not compatible with either of these drivers, you have to write your own. This Note includes all the information you need to create a MIDI driver. Note that the same drivers that work with MIDI Tools (Tool #32) also work with the MIDI Synth (Tool #35). This Note collectively refers to MIDI Tools and MIDI Synth as the "MIDI tools."

Purpose of the Driver and Description of Hardware Requirements

The Apple MIDI tools communicate to the MIDI world via a simple driver. The driver's function is managing the transmission and reception of single bytes of MIDI data between the tools and the particular MIDI hardware involved. The MIDI tools operate on the assumption that the hardware has a method of interrupting the system when a character has been received and when a character can be transmitted. Since there is quite a bit of overhead in processing MIDI data, and since MIDI data can come across a standard MIDI bus at a rate of over 3000 bytes per second, it is suggested that you provide a means for your device to buffer a few characters to reduce system overhead caused by interrupts if you are designing hardware to be used with the MIDI tools.

Format of the Driver File

The driver file is a standard OMF load file, which can be created with any of the popular Apple IIgs assemblers. The file must start with a dispatch table that contains the addresses of the standard driver routines. All driver routines must be in the same segment as the dispatch table. The dispatch table should have 13 four-byte entries, each of which contains the address of the appropriate routine minus one. Table 1 contains addresses of routines in the MIDI driver to perform specific functions.

Call	Function
Init	Called to initialize the port and prime the driver
ShutDown	Called to close the port and clean up after driver
Reset	Called at reset time by the MIDI tools
IntHandler	Called when your interrupt occurs
PollRecv	Poll input the port for data

RecvIntOn	Turns on receiver interrupts
RecvIntOff	Turns off receiver interrupts
PollXmit	Polls the transmitter to see if another character can be sent
XmitIntOn	Enables transmitter interrupts
XmitIntOff	Disables transmitter interrupts
NotImp	Currently unused
NotImp	Currently unused
NotImp	Currently unused

Table 1-MIDI Driver Function Routines

Routine Calling Conventions

All driver routines are called with full 16-bit mode enabled and should exit the same way. On entry to each routine, the accumulator contains the direct page pointer that the driver should use if it wants to use the MIDI Tools' or MIDI Synth's direct page. It is the driver's responsibility to set the direct page register and restore it on exit. All other parameters are passed on the stack and should be removed from the stack before the routine exits. The MIDI tools set aside 128 bytes of space on the passed direct page for use by the driver. They are bytes \$80-\$FF.

If you want to report an error inside of any routine (except IntHandler), set the carry flag on exit and load the accumulator with the error code. Use predefined error codes whenever possible. If you need to report a device specific error, use errors in the range \$C0-\$FF. The MIDI tools will set the high byte of the error code properly for you, so you do not need to do it yourself. Table 2 lists all of the potential predefined error codes.

Error Code	Error Definition
miToolsErr (\$2004)	The required tools were not started
miNoBufErr (\$2007)	No buffer is currently allocated
miDevNotAvail (\$2080)	Requested device is not available
miDevSlotBusy (\$2081)	Requested slot is already in use
miDevBusy (\$2082)	Requested device is already in use
miDevOverrun (\$2083)	Device overrun by incoming MIDI data
miDevNoConnect (\$2084)	No connection to MIDI
miDevReadErr (\$2085)	Framing error in received MIDI data
miDevVersion (\$2086)	ROM version is incompatible with driver
miDevIntHndlr (\$2087)	Conflicting interrupt handler installed

Table 2-Predefined Error Codes

The Driver Routines

Init

This routine is called by the MIDI tools when it wants to initialize your port and tell the driver to prepare itself for the rest of the calls. Figure 1 shows how the stack looks on entry to this call.

Figure 1-The Stack on Entry to Init

The Init routine should first test to see if the port specified by SlotFlag and SlotNum is available for use. SlotNum is the number of the slot or the port that the user has requested for use, and SlotFlag indicates whether it is a built-in port or a card in a slot. After determining that the requested device is available, you should initialize the device, allocate any memory that your driver may require (beyond what is available in the direct page), and set the proper system interrupt vector to the address passed in NewIntAddr. Before setting the vector, be sure to save the old value, as the MIDI tools expect the result from this routine to be the old address stored in the vector. On exit, the stack should contain the return address and the old vector address.

ShutDown

This routine is called when the MIDI tools want your driver to release the MIDI device and prepare to be unloaded. Figure 2 shows how the stack looks on entry to this call.

Figure 2-The Stack on Entry to ShutDown

Your routine should change the interrupt vector that you used to OldIntVector. It should then deallocate all the memory that it allocated, disable all interrupts on the device, and if needed, tell the system that you are no longer using the port in question.

Reset

This routine is called when the system has been reset by the user. Figure 3 shows how the stack looks on entry to this call.

Figure 3-The Stack on Entry to Reset

All you should do at this point is attempt to deallocate any memory you were using and disable interrupts on the device you were using.

Note: Do not set the interrupt vector to OldIntVector, instead remove the value from the stack and dispose of it.

IntHandler

The IntHandler routine is called by the MIDI tools when an interrupt occurs for the vector that you are using. The MIDI driver performs some setup then calls your routine. This routine does not have any parameters on the stack.

Once called, your IntHandler routine should test the port to see if an interrupt has occurred on your device. If your device did not cause the interrupt, you should set the carry and exit as quickly as possible, reducing the system interrupt overhead.

If your device caused the interrupt, you should test the receiver to see if any bytes of data are waiting to be read. If there is data waiting, you should load that data into the accumulator and perform a JSL to the following code:

```
InBufGlue    PEA $0400
```

PHD
RTL

This code calls the MIDI tools and tell them to accept the character in the accumulator into its input buffer. After accepting the data, control is passed back to the instruction following your JSL. If you received a byte of data and an error occurred during reception, you should load the number of the error code into the y register and perform a JSL to the following code:

```
InErrGlue    PEA $0500
              PHD
              RTL
```

Again, you will regain control right after the JSL. Once in your interrupt routine, you may perform the calls above for as much data as you like. For example, if your device has a three-byte buffer, you could call InBufGlue once for each waiting character, thus reducing your interrupt overhead and possibly preventing unneeded interrupts.

If the transmitter on your device is ready to send data, you should perform a JSL to the following code:

```
OutBufGlue   PEA $8400
              PHD
              RTL
```

This routine will return with the carry set if no data is waiting to be transmitted or the carry clear if data is available. If data is waiting, the next character to send will be in the accumulator, and you should simply send it at that time. If no more data is available, you should disable transmitter interrupts and exit. The MIDI tools will re-enable transmitter interrupts the next time it has data to send.

PollRecv

The PollRecv (Poll Receive) routine is called by the MIDI tools every now and then to see if any data might be waiting to be read. There are no parameters on the stack for this call. Your driver should test to see if any data is available and transmit it all to the MIDI tools via the InBufGlue described in the IntHandler description.

PollXmit

The PollXmit (Poll Transmit) routine is called by the MIDI tools when any data is added to the MIDI output buffer. There are no parameters on the stack for this routine. Your driver should enable transmitter interrupts, test to see if it can send any data immediately, and if it can, call OutBufGlue as described in the IntHandler description to get data to send.

XmitIntOn and RecvIntOn

These routines are called when the MIDI tools want to explicitly enable transmitter or receiver interrupts. They have no parameters on the stack and should, when called, enable transmitter interrupts for XmitIntOn and receiver interrupts for RecvIntOn.

XmitIntOff and RecvIntOff

These routine are called when the MIDI tools want to explicitly disable transmitter or receiver interrupts. They have no parameters on the stack and

should, when called, disable transmitter interrupts for XmitIntOff and receiver interrupts for RecvIntOff.

NotImp

These routines are not yet implemented, but your driver should be ready to handle a call to them. When called, they should clear the accumulator, clear the carry and perform an RTL back to the MIDI tools.

A MIDI Driver Skeleton

You can use the following sample code as a basis for a MIDI driver. It is not a complete driver in itself, and you will need to add code where comments with asterisks (***) appear for it to be functional. This example is in MPW IIgs assembler format.

```
*****
* MIDI.DRVR.Aii
*
* (C) Copyright Apple Computer, Inc. 1988
* All rights reserved.
*
* by Don Marsh & Jim Mensch
* 10/26/88
*
* This is a shell that can be used to create custom MIDI drivers for use with
* the Apple MIDI tool set. This shell is not functional, but can be used as a
* starting point for creating your own custom MIDI drivers.
*
* Files:      System Macros and equates
*
*
* Modification History:
*
* Version 1.0  Mensch
*
*      10/26/88
*
*      Create first      draft
*
*****
      Include 'E16.MIDI'
      Include 'M16.MiscTool'
      Include 'E16.MiscTool'
      Include 'M16.util'

;
; Direct page usage      Note:
; MIDI drivers may use the upper half ($80-$FF) of the MIDI direct page. When
; a MIDI driver routine is called the Accumulator will contain the direct page
; pointer for the MIDI tool set. If your driver requires more storage than
; 128 bytes, it will have to allocate them itself using the memory manager.

theuserID      equ $80          ; location to store the passed user ID
PortInUse      equ theuserID+2 ; storage for the port number in use
deref          equ PortInUse+2
Temp           equ Deref+4
EJECT
```


*

DispatchTable RECORD

*

* Description: Every MIDI Driver must start with a driver dispatch table
* that contains the entry point minus 1 of each of the
* required entry points.

*

* Inputs: None

*

* Outputs: None

*

* External Refs:

Import DRVRInit
Import DRVRShutDown
Import DRVRReset
Import DRVRIntHandler
Import DRVRPollRecv
Import DRVRRecvIntOn
Import DRVRRecvIntOff
Import DRVRPollXmit
Import DRVRXmitIntOn
Import DRVRXmitIntOff
Import DRVRNotImplemented

*

* Entry Points: None

*

DC.L DRVRInit
DC.L DRVRShutDown
DC.L DRVRReset
DC.L DRVRIntHandler
DC.L DRVRPollRecv
DC.L DRVRRecvIntOn
DC.L DRVRRecvIntOff
DC.L DRVRPollXmit
DC.L DRVRXmitIntOn
DC.L DRVRXmitIntOff
DC.L DRVRNotImplemented
DC.L DRVRNotImplemented
DC.L DRVRNotImplemented

; a few of the routines will need a temporary storage location that can be used
; even after the direct page is set back to what it was, This is a good place
; to put it!

ErrorCode ds.W 1 ; temporary holder of an error code
EndR

EJECT

*

DRVRInit PROC

*

```

* Description: This is called by the MIDI Tools when it needs to Init
* your MIDI Driver. This is usually in response to a MIDIxxx
* call made by the application.
* When this routine is called, you should allocate any buffer
* space that you will need beyond the direct page, you should
* enable the interrupts on your MIDI Device, and then set the
* appropriate system interrupt vector and return the old vector
* value. If the init works fine, clear the carry and return.
* If an error occurs return the appropriate error code
* in the Accumulator, and set the carry.
*
*

```

```

* Inputs:      UserID:Word          ID of application, for mem allocation
*             SlotFlag:Word        0 for internal port/ 1 for slot
*             SlotNum:Word         number of slot/port to use
*             NewIntVector:Long    address to give system as its new
*                                 interrupt vector. This routine is in

```

the

```

*                                 MIDI tool set, and it performs needed
*                                 setup before it calls your interrupt
*                                 routine
*

```

```

* Outputs:     OldIntVector:Long    Address interrupt vector used to have
*

```

```

* External Refs:      None
*

```

```

* Entry Points: None
*

```

```

*****
*****

```

```

; Offsets for parameters on the stack

```

```

ProcStatus      equ 1
OldDPage        equ ProcStatus+1
ReturnAddress   equ OldDPage+2
UserID          equ ReturnAddress+3
SlotFlag        equ UserID+2
SlotNum         equ SlotFlag+2
NewIntVector    equ SlotNum+2
OldIntVector    equ NewIntVector+4
ParmBytes       equ 10
ParmEnd         equ ReturnAddress+ParmBytes

```

```

; first disable interrupts since we are going to be setting up interrupt
vectors
; and enabling interrupt generating hardware. We wouldn't want an interrupt to
go
; off before we were ready to handle it! Then set us up to use the MIDI direct
; page.

```

```

        php                ; save the old proc status
        phd                ; save the old direct page
        tcd                ; Set Direct page to the one passed
        SEI                ; and disable interrupts

```

```

; now get the user ID and save it, and allocate any buffers that we may need
; Since most drivers will never need more than 128 bytes of storage we will
; not allocate any storage space

```

```

        lda UserID,s       ; first save the user ID for later
        sta theUserID      ; in our section of the MIDI DPage

```

```

; *** Insert any memory allocation needed here ***

; Next, you should check the slot flag and number to see if they are compatible
; with this driver. If they are, you should continue and initialize the proper
; port. If they are not proper, you should exit with an error.
; For this example, I will be testing the SlotFlag, to see if it is set to
; external.

```

```

        lda SlotFlag,s ; first test the slot flag to be sure
        bne FlagOK      ; its non-zero.

```

```

        ldy #miDevNotAvail ; if its zero, signal not available
        bra InitError      ; and exit via error routine

```

```

FlagOK   lda SlotNum,s ; Now save the slot number in
        sta PortInUse  ; our data area

```

```

; *** At this point you should test the firmware in the desired slot to be sure
; that the card you want is properly installed, if it is not then you should
; pass back the appropriate error ***

```

```

; Now that you know that you have the proper slot information and you have
tested
; to be sure that you have the hardware needed for the driver it is time for
you
; to initialize the interface and to enable its interrupts.

```

```

; *** Install code to initialize your hardware/interrupts here ***

```

```

; Now that the Port has been properly initialized, you must set up the proper
; system interrupt vector. Since we required an external card above it would
; make sense that you need to use the "Other unspecified interrupt handler"
; vector (Number $0017). But first, remember to get the original vector pointer
; because we must return it to the MIDI tools.

```

```

        PushLong #0 ; space for result
        PushWord #otherIntHnd ; vector to retrieve
        _GetVector ; and get the vector in question
        PullLong Temp ; place in storage for a sec

```

```

        lda Temp ; now place it on the stack
        sta OldIntVector ; as the result of this function
        lda Temp+2
        sta OldIntVector+2

```

```

        lda NewIntVector ; now move the MIDI Interrupt routine
        sta Temp ; pointer into temporary storage
        lda NewIntVector+2
        sta Temp+2

```

```

        PushWord #otherIntHnd ; now set the vector to point to
        PushLong Temp ; the MIDI drivers interrupt routine
        _SetVector

```

```

; The driver is now all set up, pull off the passed parms and we are done!
Done   ldy #0 ; set the error code to 0. No error
;

```

```

; This is the alternate label for the Done routine that should be called when
; an error has occurred.

```

```

InitError

```

```

        lda ReturnAddress,s ; Move the return address below the
        sta ParmEnd,s ; parameters

```

```

lda ReturnAddress+1,s
sta ParmEnd+1,s

pld                ; get the direct page back
plp                ; get the processor status back

tsc                ; now adjust the stack pointer
sec                ; so that the parameters are gone
sbc #ParmBytes
tcs                ; now the return address is on Top

tya                ; put any error into <A>
cmp #1             ; set the carry if non-zero
RTL                ; and return

EndP

```

EJECT

```

*****
*****

```

```

*
DRVRShutDown      PROC
*
* Description:  This routine will be called whenever the MIDI Tools want
*              to cause your driver to let go of the port it was using.
*
*
* Inputs:       OldIntVector:Long      Address to place back into the system
*                                     interrupt vector you were using
*
* Outputs:      Carry clear if successful
*               Carry set if not, error in <A>
*
* External Refs:
*               Import DrvrRecvIntOff
*               Import DRVRXMitIntOff
*
* Entry Points:
*

```

```

*****
*****

```

With DispatchTable

```

ProcStatus        equ 1
OldDPPage         equ ProcStatus+1
ReturnAddress     equ OldDPPage+2
OldIntVector      equ ReturnAddress+3
ParmBytes         equ 4
ParmEnd equ ReturnAddress+ParmBytes

```

```

; first disable interrupts since we are going to be setting up interrupt
vectors
; We wouldn't want an interrupt to go off before we were ready to handle it!
; Then set us up to use the MIDI direct page.

```

```

php                ; save the old proc status
phd                ; save the old direct page
tcd                ; Set Direct page to the one passed
SEI                ; and disable interrupts

lda #0             ; zero out the temp error code
sta >ErrorCode

```

; Now First, re-install the old interrupt vector

```
    lda OldIntVector      ; get the old vector off the stack
    sta Temp              ; and save it in globals for a sec
    lda OldIntVector+2
    sta Temp+2
```

```
    PushWord #otherIntHnd ; now set the vector to point to
    PushLong Temp          ; its original routine.
    _SetVector
```

; Next, turn off the interface hardware, and tell it to stop generating
; interrupts. We can share some code here and call our DRVRRcvIntOff and
; DRVRXmitIntOff routines. Always remember load the direct page into the
; accumulator.

```
    tdc                  ; get direct page into <A>
    js1 DRVRXmitIntOff   ; and turn off transmitter interrupts
```

```
    tdc
    js1 DRVRRcvIntOff    ; and now receiver interrupts.
```

; *** Usually turning off interrupts will be all that you would need to do at
; this point, however, if your interface card requires extra shutdown code
; this is where you would place it ***

; *** If you allocated any memory in the DRVRIInit call, this is the place to
; get rid of it.

; If an error were to occur in this routine, you should simply store the error
; number in our temporary error code variable like this

```
    ;
    ;    lda #ErrorNumber
    ;    Sta >ErrorCode
```

Done

; Now that we are done shutting down the driver, pull off the passed data
; and end.

```
    pld                  ; first retrieve the old dpage
    plp                  ; and processor status
```

```
    Longa Off            ; next move the return address
    SEP #$20             ; we need a short acc for this trick
```

```
    pla                  ; pull the 3 byte return address
    ply                  ; into <A> and <Y>
```

```
    plx                  ; now remove the remaining bytes
    plx                  ; of passed parameters
```

```
    phy                  ; and restore the return address
    pha
```

```
    Longa On
    REP #$30             ; and turn back on full 16-bit mode
```

```
    lda >ErrorCode      ; retrieve the error code
    cmp #1               ; and set the carry if non-zero
    RTL
    EndP
```

EJECT

DRVRReset PROC

* Description: This routine will be called whenever MIDIReset is called.
* and that should only happen when an actual reset occurred.
* It should in most cases perform the exact same functions
* as MIDI Shutdown.

* Inputs: OldIntVector:Long Original contents of interrupt vector

* Outputs: None

* External Refs:

* Entry Points:

jmp DRVRShutDown

EndP

EJECT

DRVRIntHandler PROC

* Description: This routine is the very core of the MIDI driver. It takes
* care of passing data back and forth between the MIDI tools
* and your hardware. It will be called for both input and
* output.

* Inputs: None

* Outputs: Carry set if interrupt not serviced

* External Refs:
Import DRVRXmitIntOff

* Entry Points:
Export InBufGlue
Export InErrGlue
Export OutBufGlue

phd ; first, save the current dpage
tcd ; and use the MIDI DPage

; The first thing the interrupt routine should do is to test to see if the
; interrupt was actually generated by our port. If it was then we should handle
; it, but if not, we should simply exit this routine with the carry set as
; fast as we can, so that the next interrupt handler will get it in a timely

```

; manner.

; *** Insert code here to test to see if the original interrupt was yours ***
        beq ServicePort ; if it was our, handle it

; If the interrupt was not ours, set the carry and leave
        pld                ; restore the direct page
        sec
        rtl

ServicePort                ; the interrupt was ours, continue

; This routine should test the interrupt again, too see if the port is ready
; to transmit or receive, If it is ready to transmit or receive, it should
; then call the ServiceRecv, or ServiceXmit routines

; *** Insert code here to test for receive

        bne ServiceRecv ; if chars waiting try receive it

; If no more characters are waiting, see if we are ready to transmit any
; characters.

        bne ServiceXmit ; if can send a character do it

; If both the above tests fail, then exit the interrupt handler for now
        pld                ; restore the direct page
        clc                ; clear the carry to indicate serviced
        RTL                ; and return

; The following routine ServiceRecv will be called when a character is waiting
; It should retrieve that character, pass it to the MIDI drivers, and then
; branch back to the beginning of ServicePort, to see if any more chars are
; waiting.
ServiceRecv

; *** Place code here that retrieves a byte of data from the port ***

; Call MIDI tools this way if no error has occurred on receive (<A> contains
the
; data read)
RecvOK
        js1 InBufGlue      ; call the MIDI tools
        bra ServicePort ; and check for more data in or out

; Call MIDI this way if a reception error has occurred (<A> contains the
; data read)
RecvErr
        ldy #miDevReadErr ; load Y with the error
        js1 InErrGlue     ; call the midi tools
        bra ServicePort

; The routine ServiceXmit will be called when the port is ready to send data.
; it will actually call the MIDI tools and get a character to send.
ServiceXmit

        js1 OutBufGlue ; call the MIDI tools for the next char
        bcs NoMoreData ; if the carry set then no data to send

; *** at this point the byte to transmit is in <A>, place your code to output
; it thru the port here ***

```

```
; Now that the data has been sent, you can either loop thru ServicePort again,  
; or you could simply end and wait for the next interrupt to send another  
; character. This sample will simply exit at this point  
bra Done ; after sending the character end.
```

```
; NoMoreData is called when the MIDI Tools said that they did not have any more  
; data to transmit, so we should turn off transmitter interrupts at this point  
; in case our device likes to keep interrupting if its empty.
```

```
NoMoreData  
phd ; push the direct page reg on the stack  
jsl DRVRXmitIntOff ; enable xmit interrupts  
  
Done  
pld ; restore the DPage  
clc ; signal the interrupt as handled  
rtl ; and get outta here!
```

```
; The routine inBufGlue should be called when you received a character from  
your
```

```
; port with no error and you want to pass it to the MIDI tools.
```

```
InBufGlue pea $0400 ; push on the long address of the  
phd ; direct page and a proc status byte  
RTL ; and jump back to the MIDI tools
```

```
; The routine inErrGlue should be called when you received a character from  
your
```

```
; port and an error has occurred. In this case, it should still be passed to  
the
```

```
; MIDI driver, as it may still be useful
```

```
inErrGlue pea $0500 ; push on the long address of the  
phd ; direct page and a proc status byte  
RTL ; and jump back to the MIDI tools
```

```
; The routine OutBufGlue should be called when you are ready to send a char  
; out your port. The MIDI tools will return with the character to send  
; in <A>. If the MIDI tools have no more characters to send then OutBufGlue  
; will return with the carry set.
```

```
OutBufGlue pea $8400 ; push on the long address of the  
phd ; direct page and a proc status byte  
RTL ; and jump back to the MIDI tools  
EndP
```

EJECT

```
*****
```

```
*****
```

```
*
```

```
DRVRPollRecv PROC
```

```
*
```

```
* Description: This routine is called by the MIDI tools when it wants to  
* pool the port for data instead of waiting for an interrupt.  
* its function is similar to that of the our interrupt handler  
* except that it only does input.
```

```
*
```

```
* Inputs: None
```

```
*
```

```
* Outputs: Carry set if interrupt not serviced
```

```
*
```

```
* External Refs:
```

```
Import InBufGlue  
Import InErrGlue
```

```
*
```

* Entry Points: None

*


```
        phd                ; first, save the current dpage
        tcd                ; and use the MIDI DPage
        php
        SEI
```

ServicePort ; the interrupt was ours, continue

```
; This routine should test the port too see if the port has any data for use
; to receive. If it does, it calls the MIDI tools and hands it off. Also note
; this routine will turn off interrupts, since we wouldn't want any stray
; receiver interrupts to spoil our fun and grab the data from us. (This is
; very important for certain types of ports which may signal that the port
; is ready and the generate an interrupt, thus leaving us in a situation where
; our interrupt routines could steal the interrupt right out from under us
before
; we fetched it, thus allowing us to possibly double post a character.
```

```
; *** Insert code here to test for received data ***
```

```
        bne ServiceRecv ; if chars waiting try receive it
```

```
; If no more data is waiting exit this routine.
```

```
        plp
        pld                ; restore the direct page
        clc                ; clear the carry no errors possible
        RTL                ; and return
```

```
; The following routine ServiceRecv will be called when a character is waiting
; It should retrieve that character, pass it to the MIDI drivers, and then
; branch back to the beginning of ServicePort, to see if any more chars are
; waiting.
```

ServiceRecv

```
; *** Place code here that retrieves a byte of data from the port ***
```

```
; Call MIDI tools this way if no error has occurred on receive (<A> contains
the
```

```
; data read)
```

RecvOK

```
        jsr InBufGlue      ; call the MIDI tools
        bra ServicePort ; and check for more data in or out
```

```
; Call MIDI this way if a reception error has occurred (<A> contains the
; data read)
```

RecvErr

```
        ldy #miDevReadErr  ; load Y with the error
        jsr InErrGlue      ; call the midi tools
        bra ServicePort
        EndP
        EJECT
```


*

DRVRPollXMit PROC

*

```

* Description: This routine is called when the MIDI tools wants to start
*              an output stream. The tool set calls this routine for
the
*              first character of data, and then this routine is
*              responsible for enabling transmitter interrupts and
sending
*              the character.

```

```

* Inputs:      None

```

```

* Outputs:     Carry set if interrupt not serviced

```

```

* External Refs:      None
                    Import OutBufGlue
                    Import DRVVRXmitIntOn

```

```

* Entry Points: None

```

```

*****
*****

```

```

                    phd                ; first, save the current dpage
                    tcd                ; and use the MIDI DPage
                    php                ; disable interrupts as we are now
going
                    SEI                ; to turn on xmitter interrupts.

```

```

; First see if the port is ready to send any data, if not simply exit

```

```

; *** Insert code here to test if output is ready ***

```

```

                    bcs Done           ; if not, then simply end

```

```

; The port is ready to accept a character for output so, call MIDI tools
; to get the next character

```

```

                    jsl OutBufGlue    ; get the next character
                    bcs Done           ; if carry set, no chars to xmit so end

```

```

                    pha                ; save the character to send
                    phd                ; push the direct page reg on the stack
                    jsl DRVVRXmitIntOn ; enable xmit interrupts
                    pla                ; retrieve the character to send

```

```

; *** Insert code here to transmit a character ***

```

```

Done

```

```

                    plp                ; get the old interrupt status
                    pld                ; get the old direct page
                    lda #0             ; no errors are possible
                    clc
                    rtl

```

```

                    EndP

```

```

                    EJECT

```

```

*****
*****

```

```

*
DRVVRXmitIntOn PROC

```

```
*
* Description: This routine will be called when the MIDI tools need to
*             enable transmitter interrupts on your device.
*
```

```
* Inputs:      None
*
```

```
* Outputs:     None
*
```

```
* External Refs:
*
```

```
* Entry Points:
*
```

```
*****
*****
```

```
        php                ; save proc status/interrupt state
        phd                ; save the old direct page
        tcd                ; use the MIDI tools DPage
        SEI                ; disable interrupts
```

```
; *** Insert code here to enable transmitter interrupts on your device
```

```
        pld                ; recover old direct page
        plp                ; recover old interrupt state
        lda #0             ; and return no-error (none possible)
        clc
        rtl
        EndP
```

```
*****
*****
```

```
*
DRVRXmitIntOff PROC
```

```
* Description: This routine will be called when the MIDI tools need to
*             Disable transmitter interrupts on your device.
*
```

```
* Inputs:      None
*
```

```
* Outputs:     None
*
```

```
* External Refs:
*
```

```
* Entry Points:
*
```

```
*****
*****
```

```
        php                ; save proc status/interrupt state
        phd                ; save the old direct page
        tcd                ; use the MIDI tools DPage
        SEI                ; disable interrupts
```

```
; *** Insert code here to Disable transmitter interrupts on your device
```

```
        pld                ; recover old direct page
        plp                ; recover old interrupt state
        lda #0             ; and return no-error (none possible)
        clc
        rtl
```

EndP

EJECT

*

DRVRRecvIntOn PROC

*

* Description: This routine will be called when the MIDI tools need to
* enable receiver interrupts on your device.

*

*

* Inputs: None

*

* Outputs: None

*

* External Refs:

*

* Entry Points:

*

php ; save proc status/interrupt state
phd ; save the old direct page
tcd ; use the MIDI tools DPage
SEI ; disable interrupts

; *** Insert code here to enable receiver interrupts on your device

pld ; recover old direct page
plp ; recover old interrupt state
lda #0 ; and return no-error (none possible)
clc
rtl
EndP

*

DRVRRecvIntOff PROC

*

* Description: This routine will be called when the MIDI tools need to
* Disable receiver interrupts on your device.

*

*

* Inputs: None

*

* Outputs: None

*

* External Refs:

*

* Entry Points:

*

php ; save proc status/interrupt state
phd ; save the old direct page
tcd ; use the MIDI tools DPage

```
SEI ; disable interrupts

; *** Insert code here to Disable receiver interrupts on your device

    pld ; recover old direct page
    plp ; recover old interrupt state
    lda #0 ; and return no-error (none possible)
    clc
    rtl
    EndP
```

```
*****
*****
```

```
*
DRVRNotImplemented PROC
```

```
*
* Description: Dummy routine, should leave the stack alone and return
* no error
```

```
*
* Inputs: None
```

```
*
* Outputs: None
```

```
*
* External Refs:
```

```
*
* Entry Points:
```

```
*****
*****
```

```
    lda #0
    clc
    RTL
    EndP
```

```
END
```

Further Reference:

- o Apple IIgs Toolbox Reference Update

Apple IIGS
#55: Avoiding ClrHeartBeat

Written by: Matt Deatherage

July 1989

This Technical Note lists changes to the description for ClrHeartBeat. This information supersedes the description in the Apple IIGS Toolbox Reference Manual.

The Apple IIGS Toolbox Reference Manual gives the following cautionary note in the description for the call ClrHeartBeat:

"A desk accessory may have installed tasks in the Heartbeat Interrupt Task queue. If you make a ClrHeartBeat call, you will remove those tasks. Therefore, under normal circumstances you should not make this call."

This isn't rude enough to get the point across to some people, so we'll try again:

The Heartbeat Interrupt Task queue does not belong to the application. Different portions of System Software can, and will, install Heartbeat Tasks. If these tasks are removed, anything from a system crash to media corruption may result. Nothing but System Software should make this call.

Further Reference

- o Apple IIGS Toolbox Reference Manual

Apple IIgs
#56: Managing Dynamic Segments

Revised by: Matt Deatherage
Written by: Eric Soldan

November 1990
July 1989

This Technical Note discusses application difficulties when transferring control to dynamic segments during low-memory conditions.

Changes since July 1989: The information formerly covered in this Note is now discussed in greater detail in Apple IIgs Technical Note #22, Proper Use of Dynamic Segments.

This Note formerly warned of the dangers of using dynamic segments-if memory is not available for the dynamic segment, the system crashes. Apple IIgs Technical Note #22, Proper Use of Dynamic Segments, covers this problem and strategies for working around it.

Apple IIGS
#57: Preventing Memory Compacting and Purging

Revised by: Dave "nocturnal" Lyons
Written by: Dave Lyons

December 1991
July 1989

This Technical Note discusses how you can use the Memory Manager from interrupt routines and documents a flag byte that debugging utilities can use to temporarily prevent the Memory Manager from moving or purging memory. Changes since July 1989: Expanded and retitled Note to discuss safe use of the Memory Manager at interrupt time.

The Memory Manager does not disable interrupts while it's busy. Instead, it increments the system BUSY flag when it's in the middle of something important.

Can interrupt routines call the Memory Manager?

If you write code that executes at interrupt time, you must check the BUSY flag (the byte at \$E100FF) before making any Memory Manager calls. If the BUSY flag is zero, it's safe to call the Memory Manager. If the BUSY flag is nonzero, the Memory Manager may be in the middle of a call, so it is not safe to call it.

What routines must check the BUSY flag?

Classic desk accessory main routines and shutdown routines do not need to check the BUSY flag. If the Event Manager is active, the CDA gets control during GetNextEvent, not at interrupt time. If the Event Manager is not active, the CDA gets control only when the BUSY flag reaches zero.

GS/OS signal handlers do not need to check the BUSY flag, because the system dispatches signals only when the BUSY flag is zero.

Run Queue tasks do not need to check the BUSY flag before calling the Memory Manager. The system dispatches Run Queue tasks at SystemTask time--the BUSY flag may not be zero, but no Memory Manager call will be in progress.

Heartbeat interrupt tasks and other interrupt handlers do need to check the BUSY flag before calling the Memory Manager.

Interrupt-time use of moveable memory blocks

If an interrupt-time routine needs access to an unlocked, non-fixed memory block, you must check the BUSY flag. It is not sufficient to lock the block, use it, and then unlock it (even if you twiddle the handle's access word directly). If the BUSY flag is non-zero, the Memory Manager could be in the middle of compacting memory, which means your block could be "in transit" from one address to another (some bytes copied, some not).

To use already-allocated memory at interrupt time, either keep the block locked or fixed, or check that the BUSY flag is zero before using the memory at interrupt time.

What if BUSY is nonzero?

If the BUSY flag is nonzero, you may want to (depending on your application) exit the interrupt routine and hope the BUSY flag is zero the next time, or call SchAddTask in the Scheduler to make the system call your routine when the BUSY flag next returns to zero. Keep in mind, though, that only four scheduled tasks can be pending at a time.

Interrupt-time flag byte

If the byte at location \$E100CB is non-zero, the Memory Manager will not move any memory blocks, and it will not purge any blocks while trying to allocate memory (PurgeHandle and PurgeAll will still purge blocks).

Debugging utilities may temporarily increment this byte to allocate memory in situations when it is not safe for existing memory blocks to be moved or purged.

This flag byte is for use only by debugging aids and System Software. It would be mind-numbingly stupid for an application to use this flag instead of using HLock and HUnlock, since the advantages of a Memory Manager architecture with relocatable blocks would be lost.

It is not useful to check the value of the \$E100CB flag. It is always set during interrupt handling whether any non-reentrant system component is busy or not.

Apple IIGS
#58: Keyboard Modifiers Register Anomaly

Written by: Dave Lyons

July 1989

This Technical Note discusses an anomaly with the keyboard modifiers register at location \$C025 which prevents it from always properly reflecting the state of the Control and Shift keys.

There are two cases where pressing the Control key turns on the Shift bit instead of the Control bit in the keyboard modifiers register:

- o An arrow key (or a Control key equivalent to an arrow key) is being held down and is repeating
- o The Space bar or Delete key is being held down and repeating with the Fast Space/Delete option selected in the Control Panel

Since the Event Manager reads the modifiers byte, desktop applications may be affected by this anomaly.

Further Reference

- o Apple IIGS Hardware Reference

Apple IIgs

#59: Do Not Create Zero-Length Text Scraps

Revised by: Dave Lyons

January 1991

Written by: Dave Lyons

July 1989

This Technical Note described a problem with zero-length text scraps.
Changed since July 1989: This Note is obsolete beginning with System Software 5.0.3. There is no longer a problem with creating a text scrap of length zero.

In System Software 5.0.3 and later, LEFromScrap no longer trashes memory if you create a text scrap (scrap type 0) with length zero.

Further Reference

- o Apple IIgs Toolbox Reference, Volume 2

Apple IIgs
#60: Menu Manager Memorabilia

Revised by: Matt Deatherage, Dave Lyons, & Tim Swihart November 1990
Written by: Dave Lyons July 1989

This Technical Note discusses the Menu Manager, specifically a few anomalies and some tips for making menus your friends.
Changes since May 1990: Noted that System Software 5.0.3 fixes a bug in NewMenuBar2.

The Menu Manager Is Your Friend

In general, this is the truth. You can do all kinds of nifty things with menus, especially in System Software 5.0 and later. However, there are a few things you should know unless you generally are fond of pain in your life.

Disabling Menus Gracefully

As documented, SetMenuFlag can be used to disable and enable entire menus. When a menu is disabled, the menu title and all items within the menu are disabled. You may pull down a disabled menu, but you may not select any item within it (unless the routine MenuGlobal has been used to allow inactive menu items to be selected).

Volume 1 of the Apple IIgs Toolbox Reference says you should call DrawMenuBar if you change the appearance of a menu title with SetMenuFlag. You can do this; this is fine. It may, however, induce dizziness if used often.

A more graceful way to dim menus is to follow SetMenuFlag with HiliteMenu. Calling HiliteMenu causes the menu title to be redrawn to reflect the current (or new) highlighting and menu flags. Using HiliteMenu instead of DrawMenuBar allows you to disable and enable menus gracefully, without noticeable flicker or threat of nasty patent infringement lawsuits from strobe light manufacturers.

"System" Bars Versus "Window" Bars

As far as the Menu Manager is concerned, there are only two kinds of menu bars. One kind is in a window and the other kind is not. The former are called "window" menu bars and the latter are generally referred to as "system" menu bars.

Most people think of the System bar as the big menu bar across the top of the screen. This is encouraged by calls like SetSysBar, which takes a menu bar handle and sets the menu bar across the top of the screen to that menu bar. Trying to rename one or the other of these two concepts at this point is probably useless; instead, this Note refers to the bar across the top of the screen as the "System" bar (with a capital S), and menu bars not in windows as "system" bars (with a lowercase s).

When you start the Menu Manager, it creates an empty System bar for you.

Before System Software 5.0, most people simply called NewMenu and InsertMenu to insert menus into that System bar. All was well in the world.

When 5.0 was released, it became very easy to create a new menu bar and all the menus within it using the NewMenuBar2 call. This avoids a lot of code, and many new people use it. The problem comes with DrawMenuBar. If you simply call NewMenuBar2 to obtain your menu bar and menus from resources, then call DrawMenuBar to make them visible, you usually get an empty menu bar. Why? The windowPtr parameter passed to NewMenuBar2 determines whether or not the new menu bar created is a system bar or a window bar-it does not force the new bar to be the System (note the capital 'S') bar. So when DrawMenuBar draws the current System bar, it hasn't changed from the empty default one created by MenuStartUp.

This is why Volume 3 of Apple IIgs Toolbox Reference recommends code similar to the following:

```
menuHandle := NewMenuBar2(refDesc,menuBarTRef,NIL);
SetSysBar(menuHandle);
SetMenuBar(NIL);           {NIL makes the System bar the current menu bar}
```

if you want your menu bar to be the one across the top of the screen.

A Bug in NewMenuBar2

NewMenuBar2 is a handy thing to have around, but it does have a problem in 5.0.2 and earlier. When the Menu Manager is done with resources, it tries to use the internal toolbox call CMReleaseResource to free them in memory. However, it passes the wrong resource ID, and CMReleaseResource calls SysFailMgr if it encounters any errors at all (such as Specified resource not found).

What NewMenuBar2 does improperly is push the high word of the resource ID onto the stack twice, instead of the high word followed by the low word. Because of the way the Resource Manager operates, CMReleaseResource returns with no error if the ID passed is NIL, but the resource is not released (another good reason not to try to use the illegal value NIL as a resource ID).

If the high word of the menu bar resource is \$0000, NewMenuBar2 passes a resource ID of NIL to CMReleaseResource, which then doesn't quite release the resource, but returns no error. The menu bar resource hangs around in memory until ResourceShutDown. It's usually fairly small, so this is no loss. It still takes up less room than menu strings, which had to stay in memory until MenuShutDown.

If the high word of the menu bar resource is not zero, the bug causes CMReleaseResource to bring down the system. When using System Software 5.0.2 or earlier, make sure all menu bar resource IDs have a high word of \$0000. System Software 5.0.3 fixes this bug.

Menu and Menu Title ID Numbers

Table 13-4 in Volume 1 of Apple IIgs Toolbox Reference gives a listing of menu and menu item ID numbers. In both lists, \$0000 and \$FFFF are "reserved for internal use" and noted that \$0000 usually indicates the first menu in the bar (or first item in the menu) and \$FFFF usually indicates the last menu in the bar (or last item in the menu). Some developers have taken this to mean that they should give their first menu an ID of \$0000 and their last one an ID of \$FFFF.

This assumption is incorrect.. The Menu Manager may change these values internally to reflect such IDs, but they must not be assigned that way by an application. Some applications that use IDs of \$0000 or \$FFFF break under System Software 5.0 and later. Note that \$0000 can be used as the insertAfter parameter to InsertMenu to insert a menu at the left of a menu bar, but \$FFFF is not a valid insertAfter value.

Desk Accessories and Menus

Some desk accessory developers would like to have their NDAs insert a menu in the System menu bar. While the menu itself can be inserted, the NDA cannot detect that a user has selected an item within that menu. The application gets the event and does not know what to do with it. NDAs that need a menu can put a menu bar in their own window. Since the mouseDown event then happens within the NDA's window, the NDA gets the event and can handle it normally. Be sure to make the NDA's menu bar the current menu bar before calling MenuSelect from within your NDA (to avoid possible conflicts between NDA menu item IDs and application menu item IDs). Restore the current menu bar to the application's menu bar before returning control to the application. Failure to do so prevents the application from finding its menus. Apple IIgs Technical Note #3, Window Information Bar Use documents how to put a menu in a window's information bar.

Documentation Error in MenuSelect

Volume 1 of Apple IIgs Toolbox Reference states that MenuSelect returns the menu ID and the item ID of the selected item in the when field of the event record. This is incorrect. MenuSelect actually returns the information in the wmTaskData field of the task record (and this, in fact, is why you pass a task record and not just an event record to MenuSelect).

Menu Strings and Bank Boundaries

NewMenu takes a pointer to a string; this string must not cross a bank boundary. If it does, a menu containing random garbage may result.

If your NewMenu strings are contained in your code segments, everything is fine-code segments cannot cross bank boundaries. Depending on your development environment, strings that are not in a code segment may or may not be allowed to cross bank boundaries. If you can find no other way to guarantee the strings do not cross a bank boundary, use NewHandle to allocate blocks with attributes \$4010 (fixed, no bank cross) and copy the strings to these blocks.

If you create menus from resources, be sure the resources have their noCrossBank attribute bits set. Note that a memory block that can cross a bank boundary usually does not, so your application may be working by accident.

Note that this restriction applies only to menu strings, not the menu templates that can be used with NewMenu2.

Return Values From GetMenuTitle and GetMItem

Starting with System Software 5.0, GetMenuTitle and GetMItem can return handles and resource IDs, not just pointers. The type of data returned depends on how the menu or item was created, so existing applications are not affected. For more information, see Apple IIgs Toolbox Reference, Volume 3, Chapter 37, "New Features of the Menu Manager."

Further Reference

- o Apple IIgs Toolbox Reference, Volumes 1 & 3
- o Apple IIgs Technical Note #3, Window Information Bar Use

Apple IIGS
#61: Window Title Handles

Written by: Dave Lyons

July 1989

This Technical Note discusses extensions to SetWTitle and GetWTitle in System Software 5.0 and later which allow handles to be used as window titles.

Prior to System Software 5.0, window titles were pointers to Pascal-style strings (with a leading length byte), but now window titles can be stored in handles, with bit 31 of titlePtr set to indicate that the parameter is actually a handle.

Once you call SetWTitle with a handle for the title parameter, the handle belongs to the Window Manager, which will dispose of it when the window is closed or retitled. You must not dispose of the handle yourself, and you must not change the data it contains.

Further Reference

- o Apple IIGS Toolbox Reference, Volume 2

Apple IIGS

#62: No Non-Solid Window Background Patterns

Written by: Dave Lyons

July 1989

This Technical Note discusses why window background patterns should always be solid; non-solid patterns are not always drawn with the expected alignment.

When the Window Manager erases part of a window's content area to its port's background pattern, it is not always aligned with already-drawn parts of the window. With a solid background pattern, this has no visible effect; however, if you try to use a grid, for example, the effect is obvious.

To simulate a non-solid background pattern, just erase the desired area to the pattern you want in your update routine. For best results, use a solid background pattern of the color most common in the pattern you really want.

For example, if you want a white grid on a black background, give the window a solid black background pattern, and use FillRect during the update routine to draw the grid. If you keep the default white background pattern, the end result will be the same, but your window content will briefly be solid white before your update routine fills it with your pattern.

Further Reference

- o Apple IIGS Toolbox Reference, Volume 2

Apple IIgs

#63: Master Color Values

Revised by: Dave Lyons
Written by: Jim Luther

May 1991
July 1989

This Technical Note documents master color values used for the Apple IIgs text, text background, and border colors.

Changes since July 1989: Added information on the standard QuickDraw II 640-mode color table and provided a 320-mode color table that produces similar colors.

Border Color Values

There are times when you may want to make parts of the IIgs Super Hi-Res screen the same color as the text, text background, and border colors. This is particularly useful when using the Apple II Video Overlay Card. Table 1 lists each color using the names from the Control Panel CDA, the color register values used for that color by the color registers, and the master color value used for that color by the Super Hi-Res screen.

Color Name	Color Register Value	Master Color Value
Black	\$0	\$0000
Deep Red	\$1	\$0D03
Dark Blue	\$2	\$0009
Purple	\$3	\$0D2D
Dark Green	\$4	\$0072
Dark Gray	\$5	\$0555
Medium Blue	\$6	\$022F
Light Blue	\$7	\$06AF
Brown	\$8	\$0850
Orange	\$9	\$0F60
Light Gray	\$A	\$0AAA
Pink	\$B	\$0F98
Light Green	\$C	\$01D0
Yellow	\$D	\$0FF0
Aquamarine	\$E	\$04F9
White	\$F	\$0FFF

Table 1-Master Color Values

The Apple IIgs Hardware Reference documents the color registers at \$C022 and \$C034, and the Apple IIgs Toolbox Reference, Volume 2 documents the master color values.

Standard 640-mode Color Table

The description of dithering on pages 16-35 and 16-36 of Apple IIgs Toolbox Reference, Volume 2 is correct, but some of the color values in Table 16-5

are incorrect. Table 2 lists the standard QuickDraw II 640-mode color table:

Color Table Offset	Color Name	Master Color Value
\$0	Black	\$0000
\$1	Red	\$0F00
\$2	Green	\$00F0
\$3	White	\$0FFF
\$4	Black	\$0000
\$5	Blue	\$000F
\$6	Yellow-green	\$0FF0
\$7	White	\$0FFF
\$8	Black	\$0000
\$9	Red	\$0F00
\$A	Green	\$00F0
\$B	White	\$0FFF
\$C	Black	\$0000
\$D	Blue	\$000F
\$E	Yellow-green	\$0FF0
\$F	White	\$0FFF

Table 2-Standard 640-mode Color Table

Table 3 shows Master Color values you can use in 320-mode to get close approximations of the sixteen standard 640-mode "solid" (really dithered) 640-mode colors.

Color Table Offset	Color Name	Master Color Value
\$0	Black	\$0000
\$1	Deep Blue	\$0008
\$2	Yellow-brown	\$0880
\$3	Gray	\$0888
\$4	Red	\$0800
\$5	Purple	\$0808
\$6	Orange	\$0F80
\$7	Pink	\$0F88
\$8	Dark Green	\$0080
\$9	Aqua	\$0088
\$A	Bright Green	\$08F0
\$B	Pale Green	\$08F8
\$C	Gray	\$0888
\$D	Periwinkle Blue	\$088F
\$E	Yellow	\$0FF8
\$F	White	\$0FFF

Table 3-Standard 640-mode Color Table

Further Reference

-
- o Apple IIgs Hardware Reference, pp. 58, 76-78
 - o Apple IIgs Toolbox Reference, Volume 2, p. 16-31

Apple IIGS

#64: Apple IIGS Installer and Installer Scripts

Revised by: Jim Luther

September 1989

Written by: Jim Luther & "Jay" Schaffer

July 1989

This Technical Note describes how the Apple IIGS Installer program executes script files and documents how to write script files for it. Note that some of the information in this Note is specific to Installer V1.10. Changes since July 1989: Changed the sourcePrefix and sourcePathname field descriptions, since sourcePrefix must not be empty if any sourcePathname fields are partial pathnames.

Introduction

The Apple IIGS Installer, a utility program that is included with Apple IIGS System Software, can be used to install System Software or applications on a given volume. "Scripts" control the Installer, and they are simply lists of files with information about where and how to install those files. The user interface of the Installer is described in the Apple IIGS System Tools Manual. This Note describes how the Installer executes scripts and how to write scripts to install your applications.

Installer Setup on Disk

Setting up the Installer on your application disk is a simple procedure.

1. Copy the Installer program to your application disk.
2. Create a subdirectory (folder) named Scripts at the same directory level as the Installer program.
3. Copy your scripts into the Scripts subdirectory.

How the Installer Processes Scripts

The Installer reads script files into memory in their entirety, parses them, strips them of all comments, compacts them, then verifies them. It then checks the scriptFlags field to see if a Caution alert should be displayed. This facility permits the script writer to force the user to read the script's help message and make a choice to either continue with file manipulations or skip the installation altogether, which is especially useful when a script installation would be inappropriate on a certain volume.

The Installer then executes the script in two passes. The first pass determines if the update can be completed by calculating the total size of the files to be deleted from the destination volume and of the files to be installed. If there is not sufficient room on the destination volume, the Installer determines the amount of additional space required to complete installation (number of blocks needed divided by two, plus one), reports this result to the user in terms of kilobytes, then terminates execution of the script. It is impossible to determine directory block requirements with complete accuracy. The Installer's space calculation algorithms are good, but they are not perfect.

If the first pass determines that there is sufficient room for the complete update, the Installer continues with the second pass, deleting and copying files in accordance with the instructions contained in the script flags. The Installer "blindly" unlocks locked files and folders, creates necessary subdirectories if they do not already exist, and replaces requested files without regard to version numbers or creation dates of existing files.

The user may terminate execution of any script (and of those which follow) by pressing the Open Apple-Period key combination. The Installer checks for key-down events between every file transfer and at the end of the first pass. If the user requests termination, the Installer warns of the possibility of leaving an unknown mix of file versions on the volume and gives the user the opportunity to continue with the installation or to terminate as requested. (See the "Error Handling" section for more details.)

Scripts are typically written with the ability to remove all of their related files from a particular volume (i.e., in case of an accidental installation); however, they do not have the ability to remove directories which contain files (even if the script installed them), and they can neither recover nor list files which were deleted during the installation process.

After processing all the instructions in a script, the Installer checks to see if additional scripts are selected, and, if they are, it executes them in the order in which they appear in the update selection window until all scripts are successfully completed. Once all selected scripts are completed, the Installer notifies the user that the installation or removal process was successful.

It is important to note several facts about script execution:

- o Each script is processed from beginning to end as if it were the only script selected.
- o If the execution of a script generates an error, or if the user terminates further processing of a script, the queue is cleared of any additional scripts waiting to be executed and control returns to the user.
- o It is possible for the Installer to execute several scripts successfully before encountering one which cannot be executed due to insufficient space on the destination volume.
- o All selected scripts use the folder that the user selects as the "Application Folder."

If a user installs or removes system files (i.e., tools, fonts, drivers, etc.) from the boot volume, it may create problems. Therefore, whenever a system level update occurs on a boot volume, the Installer disables all desk accessories and closes the Sys.Resources file. When the user quits the Installer after a system level update, it alerts the user of the need to restart the system, and the default response to this alert is to restart.

Error Handling

User Cancel Request

If the user cancels script execution any time after it has started (i.e., by pressing the Open-Apple-Period key combination), the Installer treats it as an error condition since there is likely an unknown mix of file versions on the volume. In this case, the Installer gives the user the opportunity to continue with the installation or to terminate as requested. A user-initiated cancel request is not acknowledged until the current file copy or delete request is complete. Terminating script execution also clears the queue of

other scripts waiting for execution and returns control to the user.

Non-Recoverable Errors

Some errors are simply fatal. If a directory or file is corrupted, the media is bad, or the selected script is longer than 65,535 bytes, the Installer halts execution of the script and alerts the user that a fatal error has occurred with a Stop alert box. Clicking the OK button in this alert box clears the queue of other scripts waiting for execution and returns control to the user.

Script Errors and File Not Found Errors

When the Installer detects a script error or a File Not Found error, it reports the name of the source file and destination file it was processing with the normal error message. This additional information should help script writers find the offending fileSpecification field. If the error is associated with the header, no filename is reported. This condition clears the queue of other scripts waiting for execution and returns control to the user.

Volume Not Found Errors

Volume Not Found errors produce a dialog box prompting the user to insert the missing volume. If the user clicks the OK button, the Installer attempts the file access call again, but if the user clicks the Cancel button, the Installer flags it as an error condition, clears the queue of other scripts waiting for execution, and returns control to the user.

Script File Composition

A script is simply a list of instructions for the Installer, and it can specify that files be copied from a source volume to a destination volume (or directory, when applicable) or that files be removed from a destination volume. Script files are ASCII files (file type \$04) containing printable ASCII characters (i.e., with the high-bit clear). The directory in which the Installer resides must contain a directory named Scripts, in which all script files visible to that copy of the Installer must be located. Script files may not exceed 65,535 bytes in length. Any attempt to execute a script larger than this size produces a non-recoverable error.

A script consists of a header field followed by any number of fileSpecification and comment fields. These fields are separated by tildes (~). Two consecutive tildes signal the end of the script, and any additional characters past the end of script marker are ignored. Figure 1 shows the syntax diagram for a script.

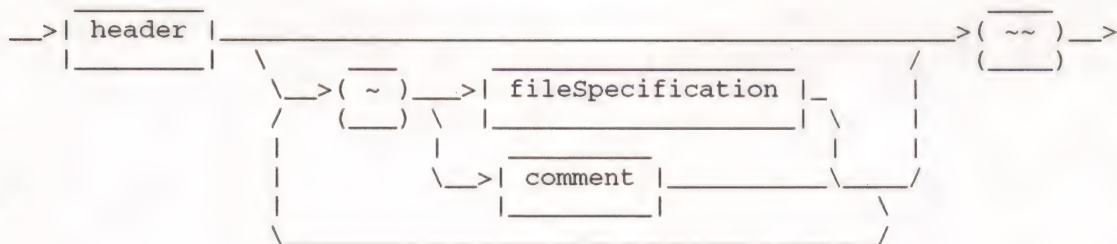


Figure 1-Script Syntax Diagram

header Field

The header field consists of the scriptIdentifier, scriptVersion, scriptFlag, scriptName, and scriptHelp fields, and it may also contain an optional sourcePrefix field. These fields supply the installer with general information about the script file. No comments are permitted within the header field. Figure 2 shows the syntax diagram for the header field.

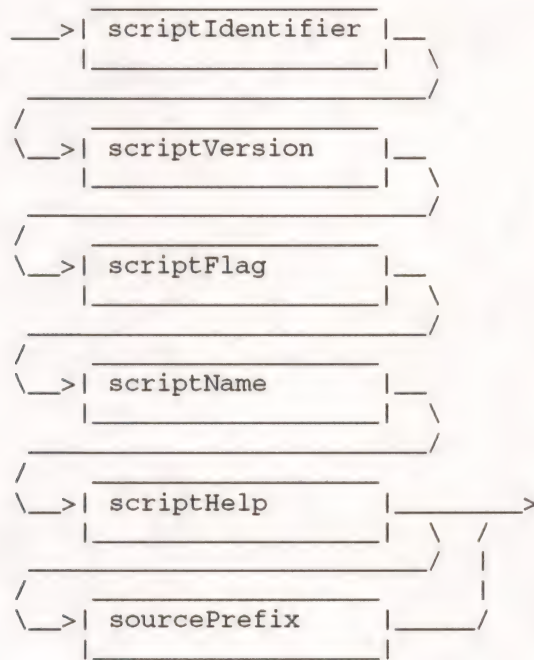


Figure 2-header Field Syntax Diagram

The scriptIdentifier field identifies the text file as a script file, and it consists of eight characters ("SCRIPT" followed by two carriage returns, or 53 43 52 49 50 54 0D 0D in hexadecimal). Figure 3 shows the syntax diagram for the scriptIdentifier field.




Figure 3-scriptIdentifier Field Syntax Diagram

The scriptVersion field defines the minimum version of the Installer program that can read and execute the instructions in this script file. It should normally consist of seven characters ("V1.10" followed by two carriage returns, or 56 31 2E 31 30 0D 0D in hexadecimal).

Version 1.0 of the Installer moves only the data fork and does not return an error. For compatibility with the original release of the Installer, the value of scriptVersion is V1.00. Scripts which move extended files (i.e., files with resource forks) or work with an AppleShare volume must have a scriptVersion of V1.10. Figure 4 shows the syntax diagram for the scriptVersion field.




Figure 4-scriptVersion Field Syntax Diagram

The scriptFlag field defines the directory requirements of the script file.

The fileSpecWorkspace field is 16 bytes that the Installer uses for work space, it can contain any character except a tilde, and it may not begin with a tilde or an asterisk (*). It is suggested that 15 readable characters followed by a carriage return might be easiest to see and count. An example of fileSpecWorkspace might be: ":::Workspace:::" followed by a carriage return, or 3A 3A 3A 57 6F 72 6B 73 70 61 63 65 3A 3A 3A 0D in hexadecimal. Figure 10 shows the syntax diagram for the fileSpecWorkspace field.




Figure 10-fileSpecWorkspace Field Syntax Diagram

The fileSpecFlags tell the Installer what this fileSpecification does. The fileSpecFlags field consists of the requiredFlag field followed by the optionalFlags field and a carriage return. Figure 11 shows the syntax diagram for the fileSpecFlags field.





Figure 11-fileSpecFlags Field Syntax Diagram

The requiredFlag field tells the Installer what to do with this fileSpecification when the Install or Remove buttons are used. The requiredFlag field must start with only one of the following characters: 1, 2, 3, or 4, and it must end with a carriage return. Any number of characters (except tilde and carriage return) may fall between the flag character and the ending carriage return. These additional characters are ignored by the Installer, making it possible to place comments within a requiredFlag field. Figure 12 shows the syntax diagram for the requiredFlag field.

The four requiredFlag characters tell the installer the following:

- 1 If the user clicks the Install button, delete the destinationPathname from the destination volume, if it exists, and copy the file from the source volume. If the user clicks the Remove button, delete the destinationPathname from the destination volume, if it exists.
- 2 If the user clicks the Install button, delete the destinationPathname from the destination volume, if it exists, and copy the file from the source volume. If the user clicks the Remove button, do nothing.
- 3 If the user clicks the Install button, delete the destinationPathname from the destination volume, if it exists. If the user clicks the Remove button, delete the destinationPathname from the destination volume, if it exists.
- 4 If the user clicks the Install button, delete the destinationPathname from the destination volume, if it exists. If the user clicks the Remove button, do nothing.




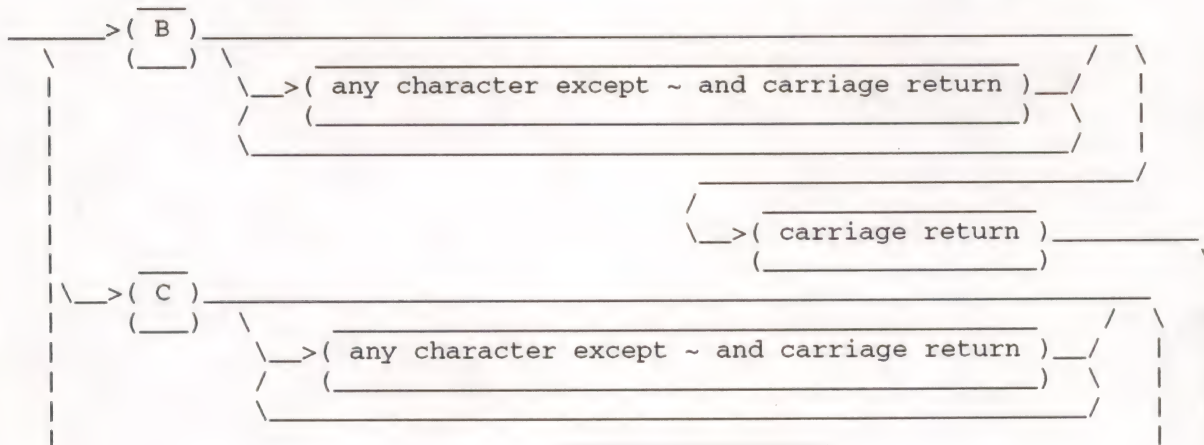


Figure 12-requiredFlag Field Syntax Diagram

The optionalFlags field gives the Installer additional duties to perform with this fileSpecification when the Install or Remove buttons are used. The five option fields, B, C, D, F, and U (must be uppercase), within the optionalFlags field are formatted the same as the requiredFlag field. Figure 13 shows the syntax diagram for the optionalFlags field.

The five optionalFlags characters tell the installer the following:

- B This flag instructs the Installer to replace the boot code on blocks zero and one of the destination volume. The boot code replacement fileSpecification is reserved for use by Apple Computer, Inc.
- C The creation date and time of the file designated by the sourcePathname field must match the createDateTime entry in this fileSpecification field.
- D The designated destinationPathname should be deleted if, and only if, it has a creation date and time that is older than createDateTime. This flag must be used with a "4" requiredFlag.
- F The file type and auxiliary type of the file designated by the sourcePathname must match the fileTypeAuxType field in this fileSpecification field.
- U Update (replace) the existing destinationPathname only if it exists. This flag must be used with a "1" or a "2" requiredFlag.



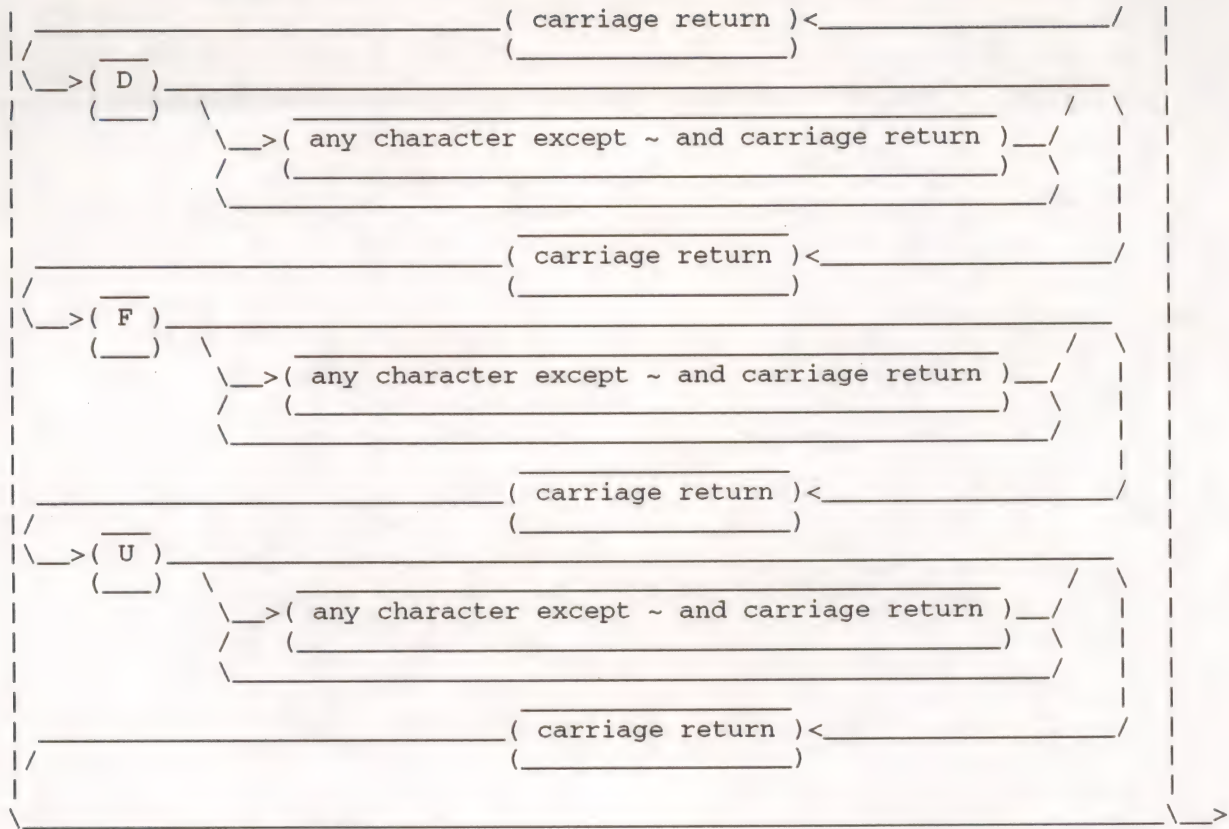


Figure 13-optionalFlags Field Syntax Diagram

The fileTypeAuxType field is used if the "F" optionalFlags field is present in the fileSpecification field. If the fileTypeAuxType field is used, it must start with a fileType field and an auxType field and must end with a carriage return. Any number of characters (except tilde and carriage return) may fall between the auxType field and the ending carriage return. These additional characters are ignored by the Installer, making it possible to place comments within the fileTypeAuxType field. If the "F" optionalFlags field is not used, then the fileTypeAuxType field must be only a carriage return. For a list of current file types and auxiliary types, see the Apple II File Type Notes. Figure 14 shows the syntax diagram for the fileTypeAuxType field.

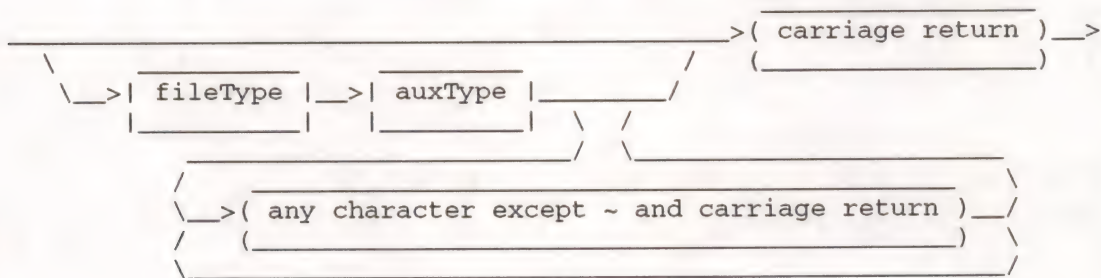


Figure 14-fileTypeAuxType Field Syntax Diagram

The fileType part of the fileTypeAuxType field consists of four, and only four, hexadecimal digits. These four digits identify a GS/OS file type if the "F" optionalFlags field is present in the fileSpecification field. An example of fileType might be: "00B3", or 30 30 42 33 in hexadecimal. Figure 15 shows the syntax diagram for the fileType field.

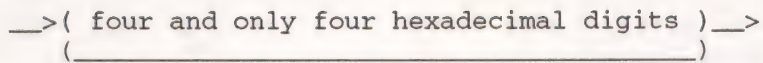


Figure 15-fileType Field Syntax Diagram

The auxType part of the fileTypeAuxType field consists of eight, and only eight, hexadecimal digits. These eight hexadecimal digits identify a GS/OS auxiliary type if the "F" optionalFlags field is present in the fileSpecification field. An example of auxType might be: "00000000", or 30 30 30 30 30 30 30 30 in hexadecimal. Figure 16 shows the syntax diagram for the auxType field.




Figure 16-auxType Field Syntax Diagram

The createDateTime field is used if the "C" or "D" optionalFlags fields are present in the fileSpecification field. If the createDateTime field is used, it must start with a date field, a single space and a time field and must end with a carriage return. Any number of characters (except tilde and carriage return) may fall between the time field and the ending carriage return. These additional characters are ignored by the Installer, making it possible to place comments within the createDateTime field. If the "C" or "D" optionalFlags fields are not used, then the createDateTime field must be only a carriage return. Figure 17 shows the syntax diagram for the createDateTime field.

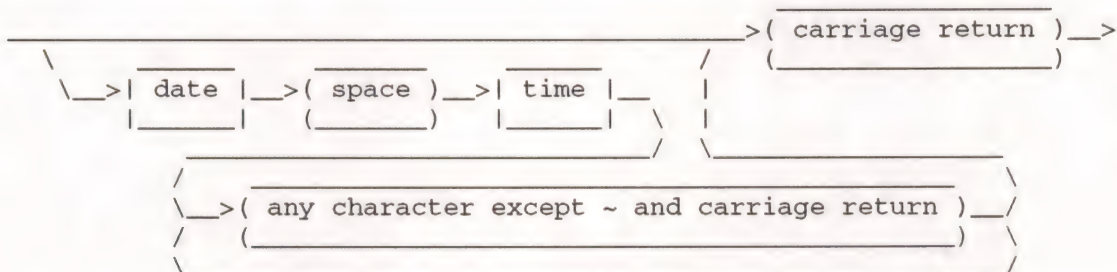
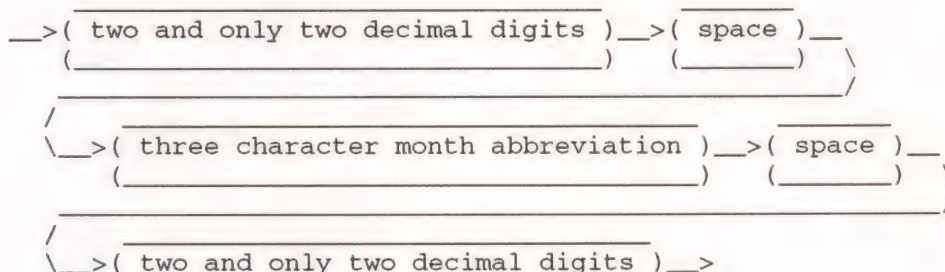


Figure 17-createDateime Field Syntax Diagram

The date subfield of the createDateime field is nine ASCII characters consisting of the day of the month, a space, a three-character month abbreviation, a space, and the year. The day of the month is a two-character number between 01 and 31. The month abbreviation may be "Jan", "Feb", "Mar", "Apr", "May", "Jun", "Jul", "Aug", "Sep", "Oct", "Nov", or "Dec" in any combination of uppercase and lowercase characters. The year is a two-character number between 00 and 99. An example of the date subfield might be: "31 Mar 57", or 33 31 20 4D 61 72 20 35 37 in hexadecimal. Figure 18 shows the syntax diagram for the date subfield.



(_____)

Figure 18-date Field Syntax Diagram

The time subfield of the createTime field is five ASCII characters consisting of the military format hour of the day, a colon, and the minute of the hour. The hour of the day is a two-character number between 00 and 23. The minute of the hour is a two-character number between 00 and 59. An example of the time subfield might be: "08:30", or 30 38 3A 33 30 in hexadecimal. Figure 19 shows the syntax diagram for the time subfield.




Figure 19-time Field Syntax Diagram

The sourcePathname field describes the name and location of the source file. The sourcePathname field consists of a valid GS/OS pathname followed by a carriage return. If the sourcePathname is a partial pathname, the sourcePrefix in the header field is used to complete the full pathname. If no sourcePrefix is specified in the header field, all sourcePathname fields must be full pathnames. If the fileSpecFlags indicate removal only, then the sourcePathname is a carriage return only. No optional comments are permitted in this field. Figure 20 shows the syntax diagram for the sourcePathname field. GS/OS Reference defines legal pathnames and prefixes.




Figure 20-sourcePathname Field Syntax Diagram

The destinationPathname field describes the name and location of the destination file. The destinationPathname field consists of a valid GS/OS partial pathname (the prefix has already been set by the Installer to the location of the destination directory, either the root directory or a user selected directory) followed by a carriage return. No optional comments are permitted in this field. Figure 21 shows the syntax diagram for the destinationPathname field. GS/OS Reference defines legal pathnames and prefixes.

Note that GS/OS now allows filenames to contain both uppercase and lowercase characters. Although filenames are not case sensitive, you should be consistent in your use of uppercase and lowercase usage in the destinationPathname field. Whatever you use here is what everyone sees.




Figure 21-destinationPathname Field Syntax Diagram

comment Field

The comment field allows commenting script files. A comment field must begin with an asterisk. The Installer ignores all characters within a comment

field, except tilde, and the comment field ends at the first tilde encountered. Figure 22 shows the syntax diagram for the comment field.

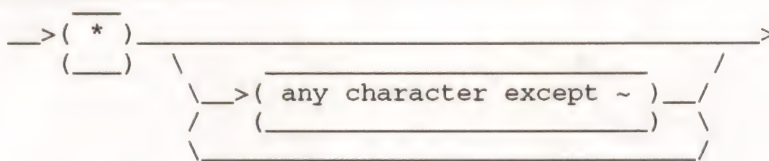


Figure 22-comment Field Syntax Diagram

Examples

Now that the script language is described, it's time to look at a couple of example scripts. The first example, CD-ROM from the System.Tools disk, installs the files necessary for you to use CD-ROM drives. The CD-ROM script is an example of using the Installer to install or update existing software. The second example, Advanced Disk Utility from the System.Tools disk, installs the files necessary to update the Advanced Disk Utility program. The Advanced Disk Utility script is an example of using the Installer to install an application in any directory on the destination volume. In both examples (Examples 1 and 2), carriage returns are shown with a paragraph mark ([p]) since they are used as delimiters within scripts.

The CD-ROM Script

The header field starts with "SCRIPT" to identify this text file as a script file. The scriptVersion is "V1.10" because this script may have to copy the resource fork of a file. The scriptFlag field is "RR", which tells the Installer to install at the root directory level and that the Remove button is valid for this script. The second "R" character in the scriptFlag field is uppercase, which tells the Installer not to display a Caution alert with the contents of the scriptHelp field. The scriptName field is "CD-ROM". The scriptName is shown in the Installer's list of scripts. The scriptHelp field (everything between the scriptName field and the "\\\" delimiter) is the text that will be displayed if the Installer's Help button is used. The sourcePrefix is ":SYSTEM.TOOLS". That is the name of the volume where the source files for this update are found.

After the header field, there is a single comment field and then five fileSpecification fields. The comment field starts at the asterisk after the first tilde and ends at the next tilde. All five fileSpecification fields start with the suggested 16-byte fileSpecWorkSpace (":::WorkSpace:::[p]") and end at the next tilde.

The first, fourth, and fifth fileSpecification fields use the "1" requiredFlag. This flag tells the Installer to copy the sourcePathname to the destinationPathname if the Install button is used, or to delete the destinationPathname if the Remove button is used. Notice the three blank lines after the "1" requiredFlag. The first blank line marks the end of the fileSpecFlags. The fileTypeAuxType field, the second blank line, is blank because the "F" optionalFlags field is not used. The createDateTime field, the third blank line, is blank because the "C" and "D" optionalFlags are not used.

The second fileSpecification field uses the "3" requiredFlag to tell the Installer to delete the destinationPathname, "System:Drivers:SCSI.Driver", if either the Install or the Delete button is used. SCSI.Driver is the interim SCSI driver from System Software 4.0. The sourcePathname field, the fourth blank line after the "3" requiredFlag, is not needed since the "3"

requiredFlag is used.

The third fileSpecification field uses the "2" requiredFlag to tell the Installer to delete the destinationPathname, "System:Drivers:SCSI.Manager" if the Install button is used. The Installer does not delete the destinationPathname if the Remove button is used. The "2" requiredFlag prevents this script from removing SCSI.Manager, which might have been installed by another script.

Two consecutive tildes after the fifth fileSpecification field signal the end of this script.

```
SCRIPT[p]
[p]
V1.10[p]
[p]
RR[p]
[p]
CD-ROM[p]
This script installs the files necessary for you to use CD-ROM drives. The
selected disk must be a startup disk.\\[p]
:SYSTEM.TOOLS~*[p]
This is the Installer script necessary to move the CD-ROM files from
:SYSTEM.TOOLS to the user's startup disk.[p]
~::~Workspace::~[~]
1[p]
[p]
[p]
[p]
System:FSTs:HS.FST[p]
System:FSTs:HS.FST[p]
~::~Workspace::~[~]
3[p]
[p]
[p]
[p]
[p]
System:Drivers:SCSI.Driver[p]
~::~Workspace::~[~]
2[p]
[p]
[p]
[p]
System:Drivers:SCSI.Manager[p]
System:Drivers:SCSI.Manager[p]
~::~Workspace::~[~]
1[p]
[p]
[p]
[p]
System:Drivers:SCSICD.Driver[p]
System:Drivers:SCSICD.Driver[p]
~::~Workspace::~[~]
1[p]
[p]
[p]
[p]
System:Desk.Accs:CDRemote[p]
System:Desk.Accs:CDRemote[p]
~~
```

Example 1-CD-ROM Script

The Advanced Disk Utility Script

The header field starts with "SCRIPT" to identify this text file as a script file. The scriptVersion is "V1.10" because this script may have to copy the resource fork of a file. The scriptFlag field is "XR", which tells the Installer the user must specify the directory where the installation should take place and that the Remove button is valid for this script. The second character (R) in the scriptFlag field is uppercase, which tells the Installer not to display a Caution alert with the contents of the scriptHelp field. The scriptName field is "Advanced Disk Utility". The scriptName will be shown in the Installer's list of scripts. The scriptHelp field (everything between the scriptName field and the "\\\" delimiter) is the text that will be displayed if the Installer's Help button is used. The sourcePrefix is ":SYSTEM.TOOLS". That is the name of the volume where the source files for this update are found.

After the header field, there is a single comment field then one fileSpecification field. The comment field starts at the asterisk after the first tilde and ends at the next tilde. The fileSpecification field starts with the suggested 16-byte fileSpecWorkspace (":::Workspace:::[p]") and ends at the next tilde.

The fileSpecification field uses the "1" requiredFlag. This tells the Installer to copy the sourcePathname to the destinationPathname if the Install button is used or to delete the destinationPathname if the Remove button is used.

Two consecutive tildes signal the end of this script.

```
SCRIPT[p]
[p]
V1.10[p]
[p]
XR[p]
[p]
Advanced Disk Utility[p]
This script installs the files necessary to update the Advanced Disk Utility
program. These files will be installed on the selected disk.\\[p]
:SYSTEM.TOOLS~*[p]
This is the Installer script necessary to update the Advanced Disk Utility
file from :SYSTEM.TOOLS to the user's disk.[p]
~:::Workspace:::[p]
1[p]
[p]
[p]
[p]
Adv.Disk.Util[p]
Adv.Disk.Util[p]
~~
```

Example 2-Advanced Disk Utility Script

Further Reference

-
- o Apple IIGS System Tools Manual
 - o GS/OS Reference

Apple IIGS

#65: Control-^ is Harder Than It Looks

Written by: Dave Lyons

September 1989

This Technical Note describes a problem using Control-^ to change the text cursor with programs that use GETLN.

On the Apple IIGS, typing Control-^ changes the cursor to the next character typed. This feature works properly from the keyboard, but there is a problem when programs print the control sequence. Try entering the following from AppleSoft to demonstrate this problem:

```
NEW
PRINT CHR$(30);"_"
```

It changes the cursor into a blinking underscore, as expected. But now enter the following:

```
12345 HOME
LIST
```

You should see 2345 HOME, which demonstrates that the first character is ignored. This is a problem with GETLN, which AppleSoft uses to read each line of input. Even if your program does not use this routine, you should be aware of this problem since it will occur the next time another program uses GETLN.

Since changing the cursor works fine when done from the keyboard, the way to work around this problem is to have your program simulate the appropriate keypresses for GETLN.

```
301: CLD                ; required by BASIC.SYSTEM
302: STA ($28),Y        ; remove cursor if present
304: LDY $0300          ; get index into simulated-keys list
307: LDA $310,Y         ; get a simulated keypress
30A: INC $0300          ; point to the next key for next time
30B: RTS                ; return the key to GETLN

310: 9E DF 8D           ; Ctrl-^, underscore, return

100 POKE 768,0:PRINT CHR$(4);"IN#A$301":REM Start getting simulated keys
110 INPUT "";A$
120 PRINT CHR$(4);"IN#0":REM Get real keys again
```

From an assembly-language program, the equivalent of IN#A\$301 is storing \$01 and \$03 in locations \$38 and \$39, while the equivalent of INPUT is JSR \$FD6A (GETLN). (Store a harmless prompt character, like \$80, into location \$33 first.)

Further Reference

- o Apple IIGS Firmware Reference, p. 77

Apple IIgs
#66: ExpressLoad Philosophy

Revised by: Matt Deatherage
Written by: Matt Deatherage

May 1992
September 1989

This Technical Note discusses the ExpressLoad feature and how it relates to the standard Loader and your application.

CHANGES SINCE SEPTEMBER 1990: Clarified some changes now that ExpressLoad and the System Loader are combined to be "Loader 4.0" in System Software 6.0. Completely removed the note about not calling Close(0) since it's not relevant.

SPEEDY THE LOADER HELPER

ExpressLoad is a GS/OS feature which is usually present with System Software 5.0 (if the ExpressLoad file is present and there's more than 512K of RAM), and always on System Software 5.0.4 and later. In fact, ExpressLoad is no longer a separate file in System Software 6.0; it's included in the System Loader version 4.0. Even though ExpressLoad is part of the Loader, we refer to its functionality separately to distinguish how the Loader takes special advantage of "expressed" files.

ExpressLoad operates on Object Module Format (OMF) files which have been "expressed," using either the APW tool Express (or it's MPW counterpart, ExpressIIgs) or created that way by a linker. Expressed files contain a dynamic data segment named either ExpressLoad or ~ExpressLoad at the beginning of the file. (Current versions of Express and ExpressIIgs create ~ExpressLoad segments, which is the preferred naming convention; older versions created ExpressLoad segments, and should be re-Expressed for future compatibility.) This segment contains information which allows the Loader to load these files more quickly, including such things as file offsets to segment headers, mappings of old segment numbers to new segment numbers (these files may have their segments rearranged for optimal performance), and file offsets to relocation dictionaries.

TWO LOADER COMPONENTS, TWO MISSIONS, ONE FUNCTION

The System Loader's function is to interpret OMF. It takes files on disk (or in memory) and transforms them from load files into relocated 65816 code. It does this very well, but in a very straightforward way. For example, when the System Loader sees the instruction to right-shift a value n times, it loads a register with the value and performs a right-shift n times.

ExpressLoad has a different mission. It relies upon the rest of the System Loader to handle OMF in a straightforward fashion so it can concentrate upon handling the most common OMF cases in the fastest possible way. For example, when asked for a specific segment in a load file, the System Loader "walks" the OMF until it finds the desired segment. ExpressLoad, however, goes directly to the desired segment since an Expressed file contains precalculated offsets to each segment in the ExpressLoad segment.

Since ExpressLoad focuses on the common things performed by the majority of

applications, it may not support those applications which rely upon certain features of OMF or the System Loader. In these cases, the System Loader loads the file as is expected.

ExpressLoad always gets first crack at loading a file, and if it is an Expressed file that ExpressLoad can handle, it loads it. If the file is not an Expressed file, the regular System Loader loads it instead. ExpressLoad also gets first shot at other loader calls.

Because an Expressed file is a standard OMF file with an additional segment, Expressed files are almost fully compatible with the System Loader (although it cannot load them any faster than before). Refer the following section for potential problems.

WORKING WITH EXPRESSLOAD

As ExpressLoad is intimate in its relationship with the System Loader, most applications work seamlessly with it; however, there are some potential problems about which you should be aware.

- o Don't mix Expressed files and normal OMF files with the same user ID. For example, if your application uses InitialLoad with a separate file, make sure that if it and your main application share the same user ID that they are both either Expressed files or normal OMF files.
- o Don't use a user ID of zero. In the past, use of zero told the System Loader to use the current user ID; however, now both the System Loader and ExpressLoad have a current user ID. Be specific about user IDs when loading. This is fixed in 6.0, but is still a good thing to avoid for compatibility with System Software 5.0 through 5.0.4.
- o Avoid loading and unloading segments by number. Since Expressed files may have their segments rearranged, if an Expressed file is loaded by the System Loader, references to segments by number may be incorrect.
- o Avoid using GetLoadSegInfo before System Software 6.0. This call returns System Loader data structures which are not supported by ExpressLoad previous to 6.0. In System Software 6.0 and later, the combined Loaders return correct information for GetLoadSegInfo regardless of whether the load file is expressed or not.
- o Don't try to load segments in files which have not been loaded with the call InitialLoad. This process was never a very good idea, and it is now apt to cause problems.
- o Don't have segments that link to other files. ExpressLoad does not support this type of link.

Further Reference

- o GS/OS Reference

Apple IIgs
#67: LaserWriter Font Mapping

Revised by: Matt Deatherage
Written by: Suki Lee & Jim Luther

May 1992
September 1989

This Technical Note discusses the methods used by the Apple IIgs Print Manager to map IIgs fonts to the PostScript(R) fonts available with an Apple LaserWriter printer.

CHANGES SINCE NOVEMBER 1989: Corrected some typographical errors and added Carta and Sonata, two fonts the LaserWriter driver knows about but aren't built into any LaserWriter.

Version 2.2 and earlier of the Apple IIgs LaserWriter driver depend solely upon font family numbers as unique font identifiers. There is a table built into the driver which maps the known font family numbers to the built-in LaserWriter family fonts. Any fonts which are not built-in are created in the printer from its bitmap font strike. Under this implementation, all font family numbers not known at the time the driver was written print using bitmap fonts. This driver knows nothing of any other fonts which may reside in the printer.

There have been many requests for the driver to take advantage of other available PostScript fonts to get high quality output from the LaserWriter. PostScript fonts from Adobe's font library, or from other PostScript font manufacturers, can be downloaded to the printer from a Macintosh and remain in the printer for use until power off. Currently there is no means to download a PostScript font with an Apple IIgs.

The Apple IIgs LaserWriter driver version 3.0 and later makes use of most resident PostScript fonts in the LaserWriter when requested. If the font is not available, then the bitmap font is used. The driver queries the printer at the start of a job for the font directory listing. The listing consists of names of all the fonts in the printer, built-in or downloaded. This information is kept locally for look up using the name of the requested font.

ISSUES

All Apple IIgs fonts contain a family name and a family number. The Apple IIgs currently identifies fonts using the family number; however, this identification method may change in the future, due to the complexity of tracking unique matches between font family names and font family numbers.

PostScript identifies its fonts by name (case sensitive) and knows nothing of any font family numbering system, Macintosh or Apple IIgs, which might be attached to a particular font. Most PostScript font families include plain, bold, italic and bold italic fonts. Some fonts families may also have serif and sans serif fonts or fonts of different weights (line thickness). These fonts are generally named by adding a style suffix to the base family name. Unfortunately, there is no uniform method for naming fonts, since most fonts were named by their designers and many of the names have historical significance.

The three examples shown in Table 1 show three variations of the plain font, two variations of the bold style, three variations of the italic style, and

three variations of the bold italic style. There are others such as ZapfChancery-MediumItalic, Korinna-KursivRegular, and LetterGothic-Slanted which all denote the italic style of the respective font family.

Style	Font names		
plain	Helvetica	Times-Roman	AvantGarde-Book
bold	Helvetica-Bold	Times-Bold	AvantGarde-Demi
italic	Helvetica-Oblique	Times-Italic	AvantGarde-BookOblique
bold italic	Helvetica-BoldOblique	Times-BoldItalic	AvantGarde-DemiOblique

Table 1-Example Font Names

The Macintosh LaserWriter driver uses a mapping scheme to compose a full PostScript font name. It relies on the Font Family Definition Record 'FOND' resource to provide a style mapping table containing the appropriate suffixes.

There are no similar resources on the Apple IIgs, which means the Apple IIgs LaserWriter driver has no way to match PostScript fonts to Apple IIgs fonts. Instead, the Apple IIgs LaserWriter driver adopts the following approach. The driver has full knowledge of all LaserWriter family built-in fonts (see Table 2 for a list of these built-in fonts) plus Carta and Sonata (two graphical fonts used in map and music programs) and uses the correct name for all style variations of the fonts. For all other fonts, the driver uses a standard set of suffixes for the style modifications. These suffixes are -Bold, -Italic, and -BoldItalic. The appropriate suffix is appended to the family name of the font, and this name is used to search the font directory table obtained from querying the printer. If a match is found, the document is printed using the corresponding PostScript font. If no match is found, then the driver tries to find the plain form of the font and creates the style modification in PostScript. A bitmap of the font is downloaded to the printer if these two searches fail.

If you are shipping your application with the intention of taking advantage of PostScript fonts when printing to a LaserWriter, please be sure to provide an Apple IIgs font whose family name is identical to the PostScript font family name.

All LaserWriters	LaserWriter Plus and LaserWriter II	
Courier	AvantGarde	Palatino
Carta	Bookman	Symbol
Helvetica	Courier	Times
Sonata	Helvetica	ZapfChancery
Symbol	Helvetica-Narrow	ZapfDingbats
Times	NewCenturySchlbk	

Table 2-Built-in LaserWriter Fonts

Further Reference

- o Apple IIgs Toolbox Reference, Volumes 1 & 2
- o Apple LaserWriter Reference

Carta is a trademark of Adobe Systems Incorporated.

PostScript and Sonata are registered trademarks of Adobe Systems Incorporated.

Helvetica(R), Palatino(R), and Times(R) are registered trademarks of Linotype Co.

ITC Avant Garde(R), ITC Bookman(R), ITC Zapf Chancery(R), and ITC Zapf
Dingbats(R) are registered trademarks of International Typeface Corporation.

Apple IIGS

#68: Tips for I/O Expansion Slot Card Design

Written by: Rob Moore & Jim Luther

September 1989

This Technical Note points out several potential problem areas developers should know about when designing I/O expansion slot cards for the Apple IIGS.

This Note is written for experienced design engineers. It is not intended to be a tutorial on Apple IIGS I/O expansion card design techniques, but rather to point out possible problem areas and pitfalls to help developers produce successful and reliable expansion cards.

The 65C816 PH2 Clock versus the Expansion Slot PH0 Clock

It is important to understand the timing of the 65C816 Phase 2 clock (PH2) on the IIGS, because several of the expansion slot signals are actually related to the PH2 clock timing, rather than the 1 MHz Phase 0 clock (PH0) available at the expansion slots. Unlike the Apple IIe, the PH2 clock at the CPU is not the same as the PH0 clock found at the expansion slots. The PH2 clock runs at a variety of periods, depending on whether the CPU is doing a normal 350 nanosecond 2.8 MHz cycle, an extended 700 nanosecond RAM refresh cycle, an isolated slow cycle, or consecutive 980 nanosecond 1.024 MHz slow cycles. During isolated slow cycles, or the first of a series of consecutive slow cycles, the fast side of the system must wait to synchronize with the 1 MHz side of the system. This synchronization results in an average cycle time of about 1.5 microseconds.

Cycle Type	Low	High	Period
Normal 2.8-MHz cycle	140ns	210ns	350ns
Refresh extended cycle	140ns	560ns	700ns
Isolated 1-MHz cycle	140ns typ.	1.33 msecs avg.	=1.5 msecs
Consecutive 1-MHz cycles	140ns	840(980)ns	980ns

Table 1-PH2 Clock Times

The Mega II Select Signal

On the Apple IIGS, the Mega II select signal (/M2SEL) is used as the enable to the slower, 1 MHz side of the system. It goes active (low) whenever the 1 MHz side RAM or I/O areas are accessed. Accesses that cause /M2SEL to be asserted include shadowed video writes, any accesses to internal I/O or expansion card slots, and accesses to banks \$E0 and \$E1. Accesses to any expansion card ROM areas that are set to Internal ROM with the Slot register do not assert the /M2SEL signal and run at the 2.8 MHz speed rather than the normal 1 MHz expansion card speed. Also, accesses to the Shadow register (\$C035), CYA register (\$C036), or DMA bank register (\$C037), and reads from the Slot register (\$C02D) or State Register (\$C068) run at full speed since they are done wholly on the fast side of the system.

/M2SEL can be viewed as an extension of the address bus on the expansion

slots. When it is active, it indicates that the CPU is running synchronized with the 1 MHz side of the system and the address on the address lines is a valid Apple II address in the 128K main or auxiliary memory space.

The Mega II Bank 0 Signal

The Mega II bank 0 signal (M2B0) provides the least significant bit of the CPU or DMA bank address to the 1 MHz side of the system. It is normally tri-stated and goes active for 140 nanoseconds, starting 140 nanoseconds after the PH0 clock falls. During the 140 nanosecond active period, M2B0 will be high whenever the CPU is accessing bank \$E1 (with the exceptions noted previously) or doing a shadowed video write or I/O access in bank \$01. Note that M2B0 does not reflect the state of the RAMRD, RAMWRT, ALTZP, 80STORE, or PAGE2 soft switches that allow access to the auxiliary 64K through bank \$00. It only indicates accesses to bank \$E1 or shadowed accesses through bank \$01.

It is generally safe to latch the state of M2B0 by using the falling edge of the Q3 clock. Even though M2B0 will be tri-stated at the about the same time as Q3 falls, the turn-off and float time on M2B0 will generally provide sufficient hold time provided that there is not more than 1 LS TTL load on M2B0.

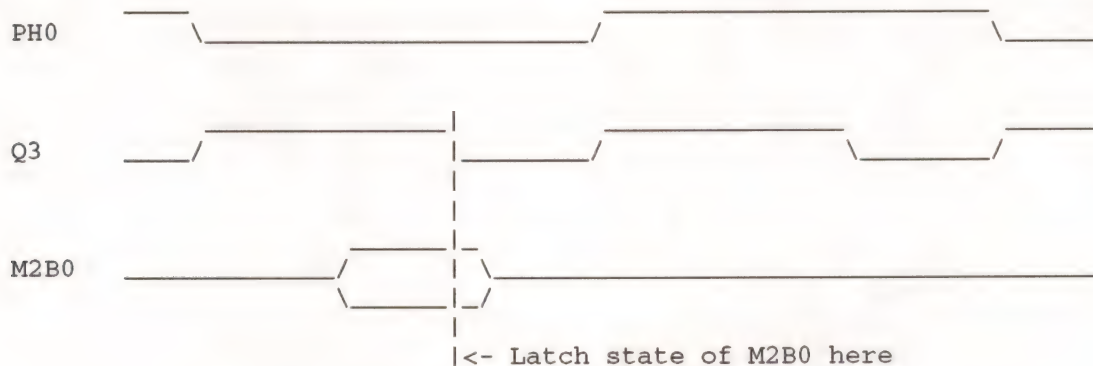


Figure 1-When to Latch State of M2B0

The Apple Video Overlay card uses M2B0 to detect writes to main and auxiliary RAM so that it can capture writes to the Apple IIGS video display buffers into its on-card display buffer. M2B0 is designed for this sort of thing and isn't of much use in most other applications. Note that M2B0 is only available on slot 3.

Using the Ready Signal

The Ready (RDY) input to the 65C816 is used to prevent a CPU cycle from completing until the expansion card has accepted the data output or has its input data available.

When the RDY input to a 65C02 or 6502 is held low, the processor continues to output the same address until RDY is released and the CPU completes the current cycle.

In the Apple IIGS, the 65C816 samples the RDY input when the PH2 clock goes low, and if RDY is low, the current CPU cycle does not complete and the address continues to be emitted. However, the bank address is not emitted while the clock is low if RDY is held low. To deal with this situation, the FPI (Fast Processor Interface) custom IC in the Apple IIGS uses a transparent latch to capture the bank address from the CPU. The latch is transparent

while the PH2 clock is low and holds the bank address while the PH2 clock is high. If RDY is low, the CPU emits an invalid bank address, so the FPI holds the latch closed while RDY is low. This action is normally completely transparent to cards in the Apple IIGS expansion slots, but if an expansion card asserts RDY while the PH2 clock is low, it is likely to cause the FPI to latch an invalid bank address, because the latch could close before the bank address from the CPU is available on the data lines.

To avoid unpredictable results, RDY should only be asserted or deasserted when /M2SEL is low and when PH0 is high, or when /DEVSEL, /IOSEL or /IOSTRB are active. When /M2SEL, /DEVSEL, /IOSEL or /IOSTRB are active, you are guaranteed that the 65C816 is running at 1 MHz and is properly synchronized to the 1 MHz side of the system. RDY should be stable at least 60 nanoseconds before the falling edge of PH0 to allow for about a 25 nanosecond skew between the PH0 slot clock and the PH2 CPU clock. Figure 2 shows where it is safe to assert or deassert RDY. Limiting changes to RDY to the time when PH0 is high guarantees that it does not change while the CPU is outputting the bank address.

The RDY line should be driven with an open-collector driver.

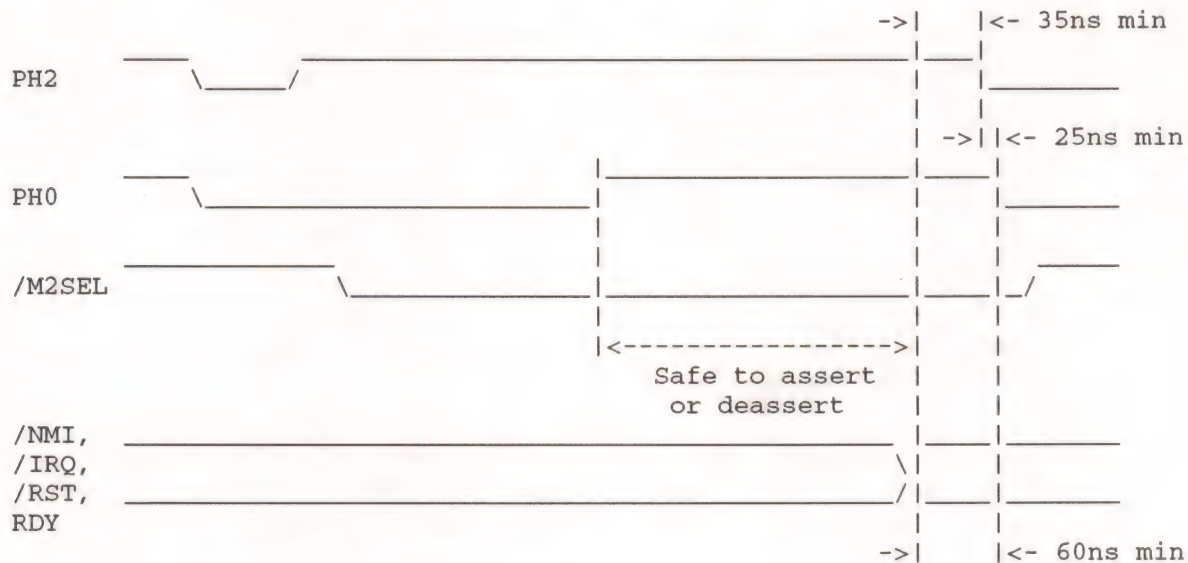


Figure 2-Control Signal Setup Time

Interrupt Request, Non-Maskable Interrupt, and Reset

The Interrupt Request (/IRQ), the Non-Maskable Interrupt (/NMI) and the Reset (/RST) signals are all interrupt lines that are sampled by the CPU when the PH2 clock falls. If they are valid 30 nanoseconds before the PH2 clock falls, they are recognized on the following cycle. If this setup time is not met, they may not be recognized until the second following cycle. Since there can be up to a 25 nanosecond skew between the PH0 and PH2 clocks, these signals should be valid 60 nanoseconds before PH0 falls if they are to be recognized on the following cycle. Figure 2 shows the correct setup time for these signals.

All three signals are all active-low and must be driven with open-collector drivers.

Note: Interrupt vectors are always pulled from ROM regardless of whether or not the language card soft-switches have ROM enabled, providing that the I/O shadowing for banks \$00/01 is enabled--which

it always is when running Apple IIGS or Apple II system software.

Direct Memory Access

The Direct Memory Access (/DMA) signal is used to temporarily halt the CPU and allow expansion cards direct access to the system RAM to transfer data at high speeds. Since the 65C816 is fully static while the PH2 clock is high (unlike the 6502), /DMA may be asserted for as long as necessary on the Apple IIGS.

The /DMA signal should be asserted and deasserted within the 100 nanosecond period after PH0 falls, and the DMA address should be emitted by the expansion card about 30 nanoseconds later. In any case, the address should be stable on the address bus no later than 120 nanoseconds after PH0 falls. This guarantees that there is enough time for the address to be decoded and for /M2SEL and M2B0 to be asserted by the FPI chip if the DMA transfer is to the 1 MHz side of the system. The bank address must be stored in the DMA bank register at location \$C037 before using DMA.

/DMA is a active-low signal and should be driven with an open-collector driver. The Apple IIGS provides a pullup for /DMA, but since the pullup is a fairly high value, it is a good idea for an expansion card that has asserted /DMA to momentarily pull it high for a few nanoseconds when deasserting it.

Note that there is a minor hardware bug in the Apple IIGS that could cause problems for developers who are unaware of it. If the CPU is currently pulling an interrupt vector when the /DMA signal is asserted, and if the DMA address is accessing the language card (\$D000-\$FFFF) space in a bank of memory where I/O and language card emulation is enabled (normally banks \$00, \$01, \$E0 and \$E1), DMA reads access ROM rather than RAM. This happens because the CPU's Vector Pull (VP) signal is active while the DMA cycle is active. Since most expansion cards that use DMA are also associated with some corresponding firmware or software driver, it's a good idea to disable interrupts prior to doing the DMA transfer, then re-enable interrupts as soon as possible after the transfer is complete. If interrupts are off too long, AppleTalk shuts down any connections to file servers because the system does not respond to AppleTalk "tickle" transactions while interrupts are disabled.

We recommend that the DMA be done with the Apple IIGS running at 1 MHz. If DMA is started during a 1 MHz cycle (/M2SEL asserted), the system continues to run slow while the /DMA signal is active.

Avoiding "Bus Fights"

The data bus on the Apple IIGS (and Apple IIe) expansion slots is a multiplexed bus that is used to carry both CPU and video display data. While PH0 is low, the bus is used to transfer data from the system RAM to the video display circuitry. When PH0 is high, the bus is available for CPU data transfers. To avoid potential (or actual) bus fights, it is helpful to avoid driving read data from an expansion card onto the bus immediately after PH0 rises. Since the video read data is driven out onto the expansion slots, and expansion card read data is driven in from the slots, it takes a finite period of time for the bus buffers to turn around. If a card drives data onto the expansion slot data bus immediately after PH0 rises, there may be a bus fight between the expansion card trying to drive the bus, and the Apple IIGS (or Apple IIe) bus buffers, which may not have turned around yet. A similar problem can occur if an expansion card leaves its read data on the bus too long after PH0 falls.

On the Apple IIGS, the data buffers turn around in 30 nanoseconds or less from the PH0 edges. Developers can avoid bus fights by simply using 74LS or 74HCT

series parts and relying upon typical delay stackups to delay driving the data bus for approximately 30 nanoseconds. A more solid technique is using the first rising edge of the 7M clock, after PH0 rises. This method may require an additional flip-flop, but it guarantees the desired delay. On the other hand, expansion card read data buffers should be turned off as soon as possible when PH0 falls to avoid a fight when the data buffers turn back out again. Figure 3 shows the recommended data transfer timing for the data bus.

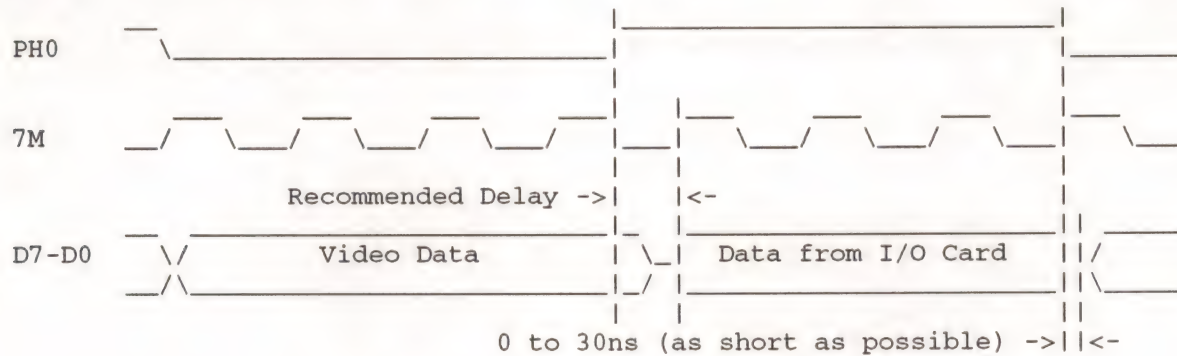


Figure 3-Recommended Data Transfer Timing

Ground Noise

Since the Apple II expansion slots were designed with only one ground pin, complex expansion cards sometimes have problems with excessive ground noise--especially in the IIGS, where the signals typically have faster rise and fall times. To reduce ground noise as much as possible, it is helpful to bypass all four supply voltages (+5 volt, +12 volt, -5 volt, -12 volt) to ground with electrolytic or solid tantalum capacitors, even if all the available voltages are not used on the expansion card. This additional bypassing has the effect of providing an improved ground by providing additional AC ground paths through the various supply pins.

To maintain a consistent ground quality over the board area on two-layer boards, it is important to properly grid the Vcc and ground traces and to fill in unused areas with ground plane.

Expansion Card Power Consumption

The Apple IIe and Apple IIGS expansion slot specifications indicate a total of 500 mA of +5 volt, 250 mA of +12 volt, 200 mA of -5 volt, and 200 mA of -12 volt power is available to all the expansion slots. With design improvements, the power required by disk drives has been reduced. Also, the Apple IIGS power supply is conservatively designed so there is somewhat more power available than indicated on the original specification. However, there is not unlimited power available, and expansion card developers should minimize power consumption as much as possible. Minimization can be accomplished by using CMOS wherever possible, using ROMs or RAMs with "power-down" mode when they are not enabled, and generally being careful to minimize parts count.

Since the Apple IIGS was released, several "super" expansion cards have become available. These cards typically provide a lot of performance and functionality, but in most cases, the power consumed by one card is more than the specified power available to all the expansion slots. Generally these cards work without problems. However, when several "super" cards are installed in a IIGS system, the total power drawn can exceed the available power supply capacity. This increase in power dissipation within the IIGS case can cause excessive heating and other associated problems when the

internal case temperatures exceed the design specifications. This could conceivably damage the IIGS power supply. Please minimize the power requirements of expansion card designs wherever possible to avoid these problems.

Further Reference

- o Apple IIGS Hardware Reference
- o Apple IIGS Firmware Reference
- o Apple IIGS Technical Note #28, Interface Card Design Guidelines
- o Apple IIGS Technical Note #32, /INH Line Anomaly

Apple IIgs

#69: The Ins and Outs of Slot Arbitration

Revised by: Matt Deatherage
Written by: Matt Deatherage

May 1990
September 1989

This Technical Note discusses the concept of a 14-slot Apple IIgs system through dynamic software slot arbitration. It presents concepts of which all IIgs programmers should be aware for full compatibility.

Changes since September 1989: Removed the section which stated that this Note showed how to switch slots in a way that does not interfere with slot arbitration and replaced it with the proper description, which is how to search a 14-slot system for peripherals and their identification bytes.

History

The Apple II has always had seven slots. In some cases (e.g., IIe), one of the slots was handled specially by the hardware, or (e.g., IIc) there was no hardware present for peripheral cards at all. But there have always been seven "slots" with firmware at location \$Cn00 (where n is the slot number). If there was no firmware, there was no peripheral connected.

With the introduction of the Apple IIgs, the Apple II family saw its first 14-slot system. Seven hardware slots are provided for peripheral cards (like on the IIe), and seven internal "ports" with connectors on the back panel are provided by the system (like on the IIc). Since \$C800 and above cannot be used for additional slots (that space is shared between all interface cards), each of the seven internal ports is matched with one of the slots, and either the port or the slot is enabled at any given time. The IIgs hardware allows switching between the two, so all fourteen slots could be used more or less simultaneously.

This situation posed a problem--the Apple II had only a disk operating system, not an overall operating system. Access to non-disk devices (i.e., character devices, like a serial card) was not arbitrated by the system in any way. The world was used to seven, and only seven, slots. Attempting to use more in a shared system such as the IIgs resulted in somebody jumping to slot firmware that somebody else had switched out. This tended to crash the system.

Then came GS/OS. With its centralized mechanism for dispatching to all devices connected to a system, GS/OS provides hope (for the first time) that a central routing mechanism can dynamically arbitrate between slots and ports, allowing the use of all 14 at one time. This is called dynamic slot arbitration, and is handled by a portion of GS/OS referred to as the Slot Arbiter.

Although the Slot Arbiter does not function in System Software 5.0 or earlier, it may function in the future. A skeleton is present in version 5.0 and later that accepts Slot Arbiter calls, but the skeleton does not actually switch any slots. This Note details the Slot Arbiter functionality and shows how to search a 14-slot system for peripherals and their identification bytes.

Note: The Slot Arbiter must not be used unless GS/OS is the current operating system.

The Slot Arbiter

The Slot Arbiter is accessed through the GS/OS system service call vector DYN_SLOT_ARBITER (\$01FCBC). On ROM 03 and later, the vector is duplicated at \$E10208. Entry to the Slot Arbiter is via a JSL instruction, and exit is via RTL. The parameters are as follows:

Entry:

A = Slot to be selected (defined below)
 X = Undefined (or Bit Encoded Slot Configuration)
 Y = Undefined
 B = Undefined
 D = Undefined
 P = N V M X D I Z C E
 x x 0 0 0 x x x 0

Exit:

A = Error Code
 X = Bit Encoded Slot Configuration
 Y = Undefined
 B = Unchanged
 D = Undefined
 P = N V M X D I Z C E
 x x 0 0 0 x x 0 0 If A = \$0000 (no error)
 x x 0 0 0 x x 1 0 If A = \$0010 (slot not available)

The slot number in the A register tells the Slot Arbiter what you are requesting. Bits 0-2 are the slot number in the range 0 through 7. Bit 3 is set if you are requesting an external slot and clear if you are requesting an internal port. Taken together, bits 0-3 give slot numbers of \$0-\$7 for internal ports and \$9-\$F for external slots. This is the same way that slot numbers are returned by the GS/OS DInfo command.

Bits 8 and 9 of the slot number indicate the action you wish the Slot Arbiter to take. A value in these two bits of 00 asks the Slot Arbiter to switch in the slot identified in bits 0 through 3. If both bits are set to 11, the Slot Arbiter restores all the slots to match the Bit Encoded Slot Configuration present in the X register. Bit Encoded Slot Configurations are discussed in the next section of this Note. Values other than 00 or 11 in bits 8 and 9 are reserved and must not be used by applications.

Bit 15 of the slot number is set if the slot selection has no slot dependencies. When the Slot Arbiter is asked to switch in a slot with no slot dependencies, it does no actual switching, although it returns a Bit Encoded Slot Configuration in the X register. The slot number and the definitions of the individual bits are illustrated in Figure 1.

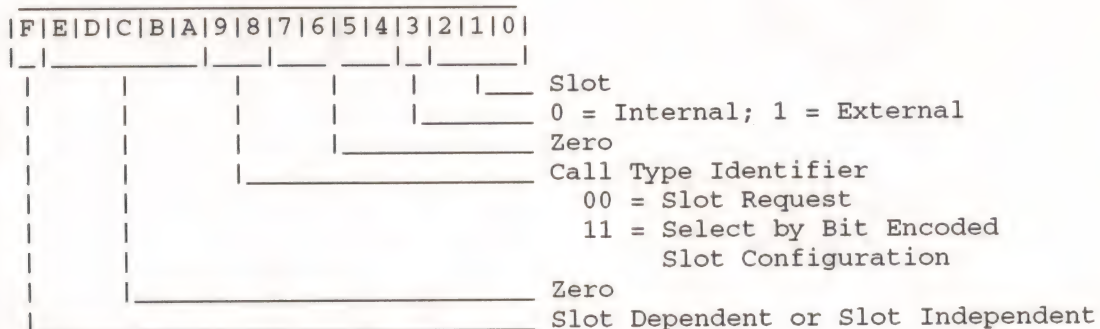


Figure 1-Slot Number and Bit Definitions

Bit Encoded Slot Configurations

Every call to the Slot Arbiter returns (on exit) a miniature picture of the slot configuration in the X register (as it was on entry). This picture has one bit set for each of the 14 slots; if the bit is set, then the corresponding slot is switched in. Bits 0 and 8 are reserved and are always clear. This picture is called a Bit Encoded Slot Configuration.

Since each external slot has the same number as an internal port (with bit 3 set), and since such pairs share the same address space, it follows that both of them may not be enabled at the same time. For example, port 5 and slot 5 (\$D) both may not be enabled. This makes the high byte of the Bit Encoded Slot Configuration the eXclusive-OR of the low byte (excluding bits 0 and 8, which are always clear). Figure 2 illustrates the Bit Encoded Slot Configuration.

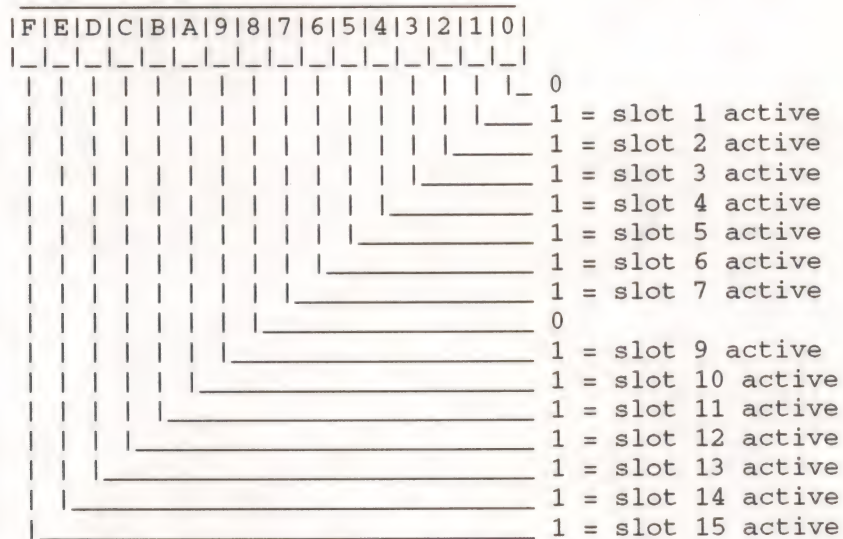


Figure 2-Bit Encoded Slot Configuration

By fully using the slot number parameter, the Slot Arbiter returns any aspect of the current slot configuration. Following are a few examples:

Slot number	Action Taken by Slot Arbiter
\$8000	Returns current Bit Encoded Slot Configuration in the X register. This number asks the Slot Arbiter to switch in with no slot dependencies (no switching), so it just returns the Bit Encoded Slot Configuration.
\$0300	Restore from Bit Encoded Slot Configuration. This command, when paired with the one above, can be used to save and restore a slot environment.
\$0005	Asks the Slot Arbiter for internal port 5.

The Impact on Applications and Drivers

Applications which correctly do all input and output through GS/OS are affected by slot arbitration, except that they find more devices available.

GS/OS uses the slot number parameter in the Device Information Block to call the Slot Arbiter, making sure the slot is available for the device before it gets control. However, there are some applications (such as peripheral card configuration programs) which go directly to firmware or hardware, not using GS/OS. Perhaps the card has no ROM, so there is no generated driver, or perhaps there is no loaded driver and the generated driver does not control certain aspects of the hardware. In any case, such applications are directly impacted by slot arbitration.

Slot Searching

The first problem is finding the hardware. In a 14-slot system, it's not suitable to just look for ID bytes between \$C100 and \$C700--two peripherals may be sharing each of those pages of slot ROM space. Drivers must examine all 14 slots, with the aid of the Slot Arbiter. The following sample code demonstrates this technique:

```

find_slot          lda    #$8000          ; request current Bit Encoded Slot
Configuration     jsr    slot_arbiter
                  phx                    ; save it on the stack

                  lda    #$000F          ; start with slot 15
                  sta    slot_number      ; be sure of the data bank when
                                          ; doing this!

slot_search        lda    slot_number      ; get the slot number to examine
                  jsr    slot_arbiter     ; and ask for it
                  bcs    continue_search ; if an error, then don't look here
                  jsr    check_for_hw     ; this routine looks for your

hardware           bcc    found_my_hw     ; if found it, we're done searching
continue_search    dec    slot_number      ; try the next lower slot
                  bpl    slot_search      ; (if there are any left, of course)

found_my_hw        plx                    ; get Bit Encoded Slot Configuration
                  ; from stack
                  lda    #$0300          ; and tell the Slot Arbiter to
                  ; restore from it
                  jsr    slot_arbiter

```

; We're done. Our slot number is in the location slot_number.

Note: You must restore the previous slot configuration when searching for a slot. This is vital to device drivers during the Drvr_Startup call, and failure to do so at other times may break older, seven-slot applications.

The Slot Arbiter attempts to maintain a static seven-slot system for applications as reflected by the user's Control Panel settings. This system allows older applications to continue to work, as something they find in an older, seven-slot scan is still present. Newer applications may wish to consider implementing a 14-slot scan, but any slot not present in the static seven-slot environment requires a call to the Slot Arbiter before and after every access to that device. The overhead in such instances may be intolerable. Apple recommends that if an application requires hardware that cannot be found in a seven-slot scan, it request the user to set the Control Panel to make the hardware available and restart the system.

Using Slot-Dependent Hardware

Applications which have slot dependencies must call the Slot Arbiter before

each use of the slot in question. Since Slot Arbitration changes the environment to which Apple IIgs programs have become accustomed, everyone has a better chance of working by sticking to the general Apple IIgs rule of "put back what you use when you're done with it." Ask for the slot, use it, then restore the previous Bit Encoded Slot Configuration. (If you use multiple slots, you might wish to get the Bit Encoded Slot Configuration, save a copy, modify it to reflect the slots you want, and restore from the modified version.)

Note: Peripherals accessed through GS/OS do not have to call the Slot Arbiter; GS/OS handles this task automatically.

There are certain applications with more specialized needs, such as high-speed, single character input or output. In such cases, the Slot Arbiter may be a bottleneck. When a slot is not switched, the Slot Arbiter returns quickly, but when a slot must be switched, it takes a significant amount of time. Doubling that significant time for switching in and restoring gives a substantial overhead for each hardware access, which may be too much for some applications.

Note: It is far better to write a GS/OS driver to deal with hardware than to write a slot-dependent application to control it. A slot-dependent application must deal with the Slot Arbiter, and the user must quit the current application to run your application just to change some aspect of the hardware. Writing a GS/OS driver lets any application, desk accessory, or CDev control your hardware with regular GS/OS calls.

Problems with Slot-Dependent Tools

Code designed before the Slot Arbiter may have slot-dependencies that cause unexpected problems when dynamic slot arbitration is fully implemented. This list includes some of the Apple IIgs System Software. Specifically, the Text Tools and the FWEntry call in the Miscellaneous Tools present problems with dynamic slot arbitration.

Text Tools

When using the Text Tools to specify a device for input, output, or error, the value specified (a four-byte parameter) is assumed to be a slot number if it is in the range 0-7. The Text Tools were not designed to use Slot Arbiter-style slot numbers, and this causes a compatibility problem.

The Text Tools were modified in System Software 5.0 to recognize Slot Arbiter-style slot numbers where possible. The trick is that it's not possible as often as we'd like. External slots are specified by using slot numbers 9 through 15; if such a slot number is used as input to a Text Tools call, the appropriate Slot Arbiter call is made and that external slot is used if it can be made available. However, internal port numbers are in the range 1-7--the same range used by the old Text Tools to indicate which of two peripherals was switched in for a particular slot. The Text Tools cannot assume that you are requesting an internal slot when using a slot number between one and seven.

For example, your old application might do a seven-slot search and find a parallel printer card in slot 1 (where the Control Panel setting for that slot is "Your Card"). If the Text Tools assumed all slot numbers in the range one through seven meant internal ports, your application would actually access the internal port 1 firmware every time it tried to access the parallel card it found in slot 1; this problem occurs since old applications don't know and don't care about internal or external slots.

The Text Tools may be used to access any external slot (if available), but

they may only be used to access internal ports that are set to internal in the Control Panel. The Text Tools slot numbers zero through seven always match the Control Panel settings.

Apple strongly recommends that the Text Tools not be used. GS/OS character-based drivers are preferable for standard character input and output. The Text Tools may be used for specialized purposes; however, you cannot access some internal ports and other components of the system that are not well-behaved. Doing so could cause your application to trash memory or media. You must assume these risks when using the Text Tools.

FWEntry

The Miscellaneous Tools call FWEntry should not be used to access entry points on a peripheral card (entry points in the \$Cxxx range). As discussed, a poorly-behaved routine could switch the slot from one you've identified to something else between the time you identify the slot and issue the FWEntry call. Furthermore, the space between \$C800 through \$CFFF cannot be identified as belonging to any given slot, and the Slot Arbiter more or less guarantees that it won't be what you expect. Accesses to peripheral card ROM space (\$Cxxx) should only be made by GS/OS drivers. FWEntry must not be used to access \$Cxxx addresses.

FWEntry is still safe to use for addresses in the \$D000-\$FFFF range.

Further Reference

-
- o Apple IIGS Toolbox Reference, Volume 2
 - o Apple IIGS Firmware Reference
 - o Apple IIGS Hardware Reference
 - o GS/OS Reference

Apple IIGS
#70: Fast Graphics Hints

Written by: Don Marsh & Jim Luther September 1989

This Technical Note discusses techniques for fast animation on the Apple IIGS.

QuickDraw II gives programmers a very generalized way to draw something to the Super Hi-Res screen or to other parts of Apple IIGS memory. Unfortunately, the overhead in QuickDraw II makes it an unacceptable tool for all but simple animations. If you bypass QuickDraw II, your application has to write pixel data directly to the Super Hi-Res graphics display buffer. It also has to control the New-Video register at \$C029, and set up the scan-line control bytes and color palettes in the graphics display buffer. Chapter 4 of the Apple IIGS Hardware Reference documents where you can find the graphics display buffer in memory and how the scan-line control bytes, color palettes, and pixel data bytes are used in Super Hi-Res graphics mode. The techniques described in this Note should be used with discretion--we do not recommend bypassing the Apple IIGS Toolbox unless it is absolutely necessary.

Map the Stack Onto Video Memory

To achieve the fastest screen updates possible, you must remove all unnecessary overhead from the instructions that perform graphics memory writes. The obvious method for achieving sequential writes to the graphics memory uses an index register, which must be incremented or decremented between writes. These operations can be avoided by using the stack. Each time a byte or word is pushed onto the stack, the stack pointer is automatically decremented by the appropriate amount. This is faster than doing an indexed store followed by a decrement instruction.

But how is the stack mapped onto the graphics memory? The stack can be located in bank \$01 instead of bank \$00 by writing to the WrCardRAM auxiliary-memory select switch at \$C005. Bank \$01 is shadowed into \$E1 by clearing bit 3 of the Shadow register at \$C035. Under these conditions, if the stack pointer is set to \$3000, the next byte pushed onto the stack is written to \$013000, then shadowed into \$E13000. The stack pointer is automatically decremented so the stage is set for another byte to be written at \$E12FFF.

Warning: While the stack is mapped into bank \$01, you may not call any firmware, toolbox or operating system routines (ProDOS 8 or GS/OS). Don't even think about it.

Unroll All Loops

Another source of overhead is branching instructions in loops. By "straight-lining" the code to move up a scan-line's worth of memory at one time, branch instructions are avoided. Following is an example of this technique.

```
lda    |164,y          ; accumulator is 16 bits for
pha                    ; best efficiency
lda    |162,y
pha
lda    |160,y
```

pha

In this example, the Y register is used to point to data to be moved to the graphics memory, and hard-coded offsets from the Y register are used to avoid register operations between writes.

Hard-Code Instructions and Data

In desperate circumstances, it is necessary to remove overhead from the previous code example. This can be accomplished by hard-coding pixel data into your code instead of loading pixel values from a separate data space and transferring them to the graphics memory (as in the example). If you are writing an arbitrary pattern of three or fewer constant values to the screen, for example, the following method is the fastest known:

```
lda    #val1
ldx    #val2
ldy    #val3
pha                    ; arbitrary pattern of pushes
phx
phy
phy
phx
```

In cases where many different values must be written to the screen, pixel data can be written to the screen using immediate push instructions:

```
pea    $5389           ; some arbitrary pixel values
pea    $2378
pea    $A3C1
pea    $39AF
```

Your program can generate this mixture of PEA instructions and pixel data itself, or it could load pixel data that already has PEA instructions intermixed (thus increasing the data size by one half).

Be Aware of Slow-Side and Fast-Side Synchronization

Estimating execution speed by counting instruction cycles is always a challenging task on the IIGS, but it is particularly tricky when one is writing to the graphics memory. The graphics memory resides in the side of the IIGS system controlled by the 1 MHz Mega II chip, which means that during all writes to this memory, the fast side of the system controlled by the Fast Processor Interface (FPI) chip must be synchronized with slow side of the system controlled by the Mega II, even if the system is running code at full native speed. This synchronization is performed automatically and transparently by the FPI in the IIGS, and it isn't normally of concern to the programmer. Animation programmers must worry about synchronization delays, however, because slight changes in graphics update code may change the frequency of these delays, and hence the speed of the program. In practical terms, this means that one loop writing data to the graphics memory may run at the same speed as a second loop with a higher cycle count.

A careful analysis of the synchronization problem leads to the following tables, which are useful as a rough estimate of the speed attained by different pieces of code. Each entry is based on the number of cycles consumed during consecutive write instructions. For example, a series of PEA instructions requires five cycles for each 16-bit write. A short PHA instruction followed by a branch requires six cycles for each 8-bit write.

Fast Cycles per Write (byte)	Actual Speed (microseconds/byte)
------------------------------	----------------------------------

3 to 5	2.0
6 to 8	3.0
9 to 11	4.0

Fast Cycles per Write (word)	Actual Speed (microseconds/word)
4 to 6	3.0
7 to 8	4.0
9 to 11	5.0

The times given in the tables apply only if the same number of fast cycles separate each consecutive write operation. The first write operation in a set of write instructions usually takes longer than subsequent writes, because the potentially long synchronization operation is accomplished at that time. Unpredictable delays caused by memory refresh slow things down further, although refresh delays byte-wide writes more often than word-wide writes. Therefore, it is usually preferable from a speed standpoint to use word-wide writes to the graphics memory.

For more information on synchronization cycle timing within the IIGS, see Chapter 2 of the Apple IIGS Hardware Reference and Apple IIGS Technical Note #68, Tips for I/O Expansion Slot Card Design.

Use Change Lists

The timing data given in the preceding section shows that it is not possible to perform full-screen updates in the time it takes the IIGS to scan the entire screen. In fact, it would be difficult to update more than one-sixth of the screen in one scan time. Therefore, it is necessary to update only those pixels which have actually changed from the previous frame of animation. One method of doing this is to precalculate the pixels which change by comparing each frame against the preceding frame. For interactive animation, fast methods must be developed for predicting which areas of the screen must be updated (a determination of the exact pixels might require more computation than the actual update would require).

Using the Video Counters

To achieve "tear-free" screen updates, it is necessary to monitor the location of the scan-line beam when writing to graphics memory. As described in Apple IIGS Technical Note #39, Mega II Video Counters, the VertCnt and HorizCnt Mega II video counter registers at \$C02E-C02F allow you to determine which scan line is currently being drawn.

By using only the VertCnt register and ignoring the low bit of the 9-bit vertical counter stored in HorizCnt, you can determine within 2 scan lines which scan line is currently being drawn. The VertCnt video counter contains the number of the current scan line divided by two, offset by \$80. For example, if the scan-line beam was currently refreshing either scan line four or five, VertCnt would contain \$82 ($4/2 + \80 or $5/2 + \$80$). Vertical blanking happens during VertCnt values \$7D through \$7F and \$E4 through \$FF.

Clever updates can modify twice as many pixels on the screen by sacrificing some smoothness, running at 30 frames per second instead of 60. The technique is as follows:

1. Wait for the scan line beam to reach the first scan line.
2. Start updates from the top of the screen, being careful not to pass the scan line beam.
3. Continue updates while the scan line beam progresses toward the

bottom of the screen, then goes into vertical blanking, then restarts at the top of the screen.

4. Finish the update before the scan line beam catches the update point.

Careful use of this method allows a frame to be updated during two scans of the screen instead of just one. If you are not sufficiently careful, tearing results.

Note: The Apple IIGS main logic board Mega II-VGC registers and interrupts are not synchronous to the Apple II Video Overlay Card video and therefore should not be used for time synchronization with the Apple II Video Overlay Card video output. However, they can be used for time synchronization with the Apple IIGS video output. See the Apple II Video Overlay Card Development Kit for more information.

Interrupts

It is not possible to support interrupts while sustaining a high graphics update rate, unless jerkiness or tearing is acceptable. Be aware that many system activities such as GS/OS and AppleTalk depend on interrupts and do not function if interrupts are disabled.

Further Reference

-
- o Apple IIGS Firmware Reference
 - o Apple IIGS Hardware Reference
 - o Apple II Video Overlay Card Development Kit
 - o Apple IIGS Technical Note #39, Mega II Video Counters
 - o Apple IIGS Technical Note #40, VBL Signal
 - o Apple IIGS Technical Note #68, Tips for I/O Expansion Slot Card Design

Apple IIgs
#71: DA Tips and Techniques

Revised by: Dave "Mr. Tangent" Lyons
Written by: Dave Lyons

May 1992
November 1989

This Technical Note presents tips and techniques for writing Desk Accessories.

CHANGES SINCE DECEMBER 1991: Reworked discussion of NDAs and Command-keystrokes. Marked obsolete steps in "NDAs Can Have Resource Forks."

CLASSIC DESK ACCESSORY TIPS AND TECHNIQUES

READING THE KEYBOARD

For a CDA that runs only under GS/OS, the Console Driver is the best choice for reading from the keyboard. Other CDAs have two cases to deal with: the Event Manager may or may not be started. The Text Tools can read the keyboard in either case, but you should avoid using the Text Tools whenever possible (see Apple IIgs Technical Note #69, The Ins and Outs of Slot Arbitration).

You can call EMStatus to determine whether the Event Manager is started. When it is, you can read keypresses by calling GetNextEvent. When the Event Manager is not started, you can read keys directly from the keyboard hardware by waiting for bit 7 of location \$E0C000 to turn on. When it does, the lower seven bits represent the key pressed. Once you've detected a keypress, you need to write to location \$E0C010 to remove the keypress from the buffer.

Alternately, you can use IntSource (in the Miscellaneous Tools) to temporarily disable keyboard interrupts and then read the keyboard hardware directly. Be sure to reactivate keyboard interrupts if, and only if, they were previously enabled.

JUST ONE PAGE OF STACK SPACE

CDAs normally have only a single page of stack space available to them (256 bytes at \$00/01xx). Your CDA may or may not be able to allocate additional stack space from bank 0 during execution. The following code (written for the MPW IIgs cross-assembler) shows a safe way to try to allocate more stack space and to switch between stacks when the space is available.

If ProDOS 8 is active, your CDA cannot allocate additional space (and there is no completely safe way to "borrow" bank 0 space from the ProDOS 8 application).

```
HowMuchStack  gequ    $1000          ;try for 4K of stack space

start          phd
               phb
               phk
               plb
               pha          ;Space for result
               pha
               PushLong #HowMuchStack
               pha
               _MMStartUp
```



```

RealCDAentry  bcs    smallStack          ;if c set, only 1 page of stack
                                           ;is available
                                           ; put something interesting here
               ...
               rti

smallStack    _SysBeep
               rti

```

Note that interrupts are disabled while the page-one stack is being restored; they are reenabled (if they were originally enabled) only after the stack pointer is safely back in page one.

INTERRUPTS, EVENT MANAGER, MEMORY, AND CDAS

Whether the Event Manager is active or not, the user hits Apple-Ctrl-Esc and usually gets to the CDA menu. It looks the same, but what happens internally is different affects what happens when your CDA allocates memory.

When the Event Manager is active (as it normally is while the user is running a Desktop application), hitting Apple-Ctrl-Esc posts a deskAcc event to the event queue. The CDA menu appears only when the application calls GetNextEvent or EventAvail with the deskAcc bit enabled in the event mask.

So with the Event Manager active, the CDA menu and individual CDAs are running in the "foreground"--no processor interrupt is being serviced, and the foreground application is stuck inside the GetNextEvent or EventAvail call. The Memory Manager knows that no interrupt is in progress, so it will happily compact and purge memory if necessary to carry out a memory allocation request from your CDA. This is just fine, since the foreground application made a toolbox call--unlocked memory blocks are not guaranteed to stay put.

When the Event Manager is not active, hitting Apple-Ctrl-Esc either enters the CDA menu immediately (if the system Busy Flag is zero) or calls SchAddTask so that the CDA menu appears during a the next DECBUSYFLG call that brings the system Busy Flag down to zero. If the CDA menu appears during a DECBUSYFLG , normal memory compaction and purging are possible, just like when the Event Manager is active.

But if the Busy Flag was zero when the user hit Apple-Ctrl-Esc, then the CDA menu appears inside of the interrupt, and the foreground application is at an unknown point where it may justifiably expect that unlocked memory blocks will not move or be purged (see Apple IIgs Toolbox Reference, Volume 1, page 12-5). (Note that the Desk Manager does a tricky dance to allow additional interrupts to occur, even though the Apple-Ctrl-Esc interrupt will not return until the user chooses Quit from the CDA menu. Normally interrupts cannot be nested; the Desk Manager and AppleTalk are exceptions.)

The Memory Manager knows an interrupt is in progress, so CompactMem takes no action and memory allocation requests do not cause unlocked memory blocks to move and do not attempt to purge purgeable blocks to make room. Memory allocation requests will still normally succeed, but you will not be able to allocate a block larger than the value returned by MaxBlock.
New Desk Accessory Tips and Techniques

AN NDA CAN FIND ITS MENU ITEM ID

After the application has called FixAppleMenu, an NDA can look at its menu item template (after the "\H" in the NDA header) to determine the menu ID corresponding to the NDA's name in the Apple menu. This is sometimes useful to pass to OpenNDA (if the NDA has some way to open itself), or to pass to a Menu Manager call.

Finding the menu item ID in the NDA's header is easy if the NDA is written in assembly. In a high-level language it may be harder (if you don't have direct access to your NDA's header, you need to find it on the fly and scan for the "\H").

NDAS AND COMMAND- KEYSTROKES

To give the user a consistent way to close NDA windows, System 6.0 handles Command-W automatically when a system window is in front. It calls CloseNDAByWinPtr without letting the NDA or the application see the Command-W.

However, there is a special action code (optionalCloseAction) that an NDA can accept to handle the Close request itself. This way the NDA can offer the user a chance to cancel the Close, which is impossible when the system calls the NDA's main Close routine, as CloseNDAByWinPtr does. (See the System 6.0 Toolbox documentation for details.)

There is no way for an NDA to accept some keystrokes and pass others along to applications, but if your NDA does not want any keystroke events, turn off the corresponding eventMask bits in the NDA header (this allows the application to receive keystrokes while your NDA window is in front).

CALLING INSTALLNDA FROM WITHIN AN NDA

It is possible to write an NDA that installs other NDAs. However, with System Software 5.0 and later, InstallNDA returns an error when called from an NDA. When your NDA has control because the Desk Manager called one of your NDA's entry points, the Desk Manager's data structures are already in use, so InstallNDA is unable to modify them.

The solution is to use SchAddTask in the Scheduler to postpone the InstallNDA call until the system is not busy. Remember that the Bank and Direct Page registers are not defined when your scheduled task is executed.

PROCESSING MOUSEUP EVENTS

When an NDA's action routine receives a mouseUp event, it is not always safe for the NDA to draw in its window.

For example, when the user drags an NDA window, the NDA receives the mouseUp before the window is actually moved, and before DragWindow erases the outline of the new window position, which may overlap the window's content. In addition, when the user chooses a menu item, the front NDA receives the mouseUp before the menu's image is removed, and the image may overlap the NDA's window. In either case, drawing in the window makes a mess.

The solution is to avoid drawing in direct response to a mouseUp. Instead, invalidate part of the window to force an update event to happen later. NDAs Can Have Resource Forks

Following is the recommended way for a New Desk Accessory to use its file's resource fork.

In the NDA's Open routine, do the following. Steps that are obsolete (and safely omitted) with System Software 6.0 and later are marked with an asterisk (*):

1. Call GetCurResourceApp and keep the result.
2. If the NDA does not already know its Memory Manager user ID, call MMStartUp to get it.
3. Call ResourceStartUp using the NDA's user ID.

4. Call the Loader function LGetPathname2 with the NDA's user ID (and a fileNumber of \$0001) to get a pointer to the NDA's pathname. (The result is a pointer to a class-one GS/OS string.)
- *5. Use GetLevel to get the current file level, then use SetLevel to set it to zero. This helps protect your resource fork from being closed accidentally.
6. Use GetSysPrefs to get the current OS preferences, then use SetSysPrefs to ensure that the user is prompted, if necessary, to insert the disk containing your resource fork. (To compute the new preferences word, take the current one, AND it with \$1FFF, and ORA it with \$8000. This tells GS/OS to deal with volume-not-found conditions by putting up a please-insert-disk dialog with an OK button and a Cancel button.)
7. Call OpenResourceFile using the result from LGetPathname2. Save the returned fileID--you need it when closing the file. (Be prepared to deal with an error, such as \$0045, Volume Not Found.)
8. Use SetSysPrefs to restore the OS preferences saved in step six.
- *9. Use SetLevel to restore the file level to its old value (saved in step five).
10. Call SetCurResourceApp with the old value saved in step one.

In the NDA's action routine, no special calls are necessary--the Desk Manager calls SetCurResourceApp automatically before calling your action routine, so your NDA's own resource search path is already in effect.

Run queue routines and NDA installs with AddToRunQ are treated the same way--the NDA's resource search path is automatically in effect when the run queue routine is called.

In the NDA's Close routine, do the following:

1. Call CloseResourceFile with the fileID that was returned when you opened it.
2. Call ResourceShutDown with no parameters.

NDAS MUST BE CAREFUL HANDLING MODAL WINDOWS

If your NDA uses its resource fork and calls TaskMaster with a restricted wmTaskMask to produce a modal window, you must be careful not to allow TaskMaster to update the contents of any application windows that happen to need updating.

The problem is that an application window's wContDraw routine can reasonably assume that the current Resource Manager search path is the application's, but TaskMaster does not take any special steps to set it. When the content-draw routine draws controls which were created from resources which are not presently in the resource search path, the system may crash.

If your NDA does not start up the Resource Manager, the Desk Manager is unable to SetCurResourceApp to your NDA, so the application's search path is still in effect--no problem. But if your NDA does start the Resource Manager, you have to be careful not to cause application routines to be called.

AVOID HARD-CODING YOUR PATHNAME

If your NDA needs to know its own pathname or the pathname of the directory it's in, call LGetPathname or LGetPathname2 using your User ID. This is a better method than hard-coding "*:System:Desk.Accs:MyDAName" because the user may change your DA's file name or use a utility to install it from some non-standard directory.

AVOID EXTRA GETNEWID CALLS

Normally there is no reason for a Desk Accessory to call GetNewID. When you can, just call MMStartUp to find your own User ID, and use that. You can freely use all the auxiliary IDs derived from your main ID (MMStartUp+\$0100, MMStartUp+\$0200, ..., MMStartUp+\$0F00).

By not calling GetNewID, you conserve the limited supply of IDs (255 of in the \$50xx range for Desk Accessories), and you make life easier for people trying to debug their systems, since all your allocated memory can be readily identified.

OPEN IS NOT CALLED IF NDA IS ALREADY OPEN

Your NDA's Open routine does not get called if the user chooses the NDA from the Apple menu while the NDA is already open. In this case, the Desk Manager simply calls SelectWindow on your existing window.

There is no need to include code in your Open routine to check if your window is already open, and to call SelectWindow if it is.

Further Reference

- o Apple IIgs Toolbox Reference, Volumes 1-3
- o GS/OS Reference
- o Apple IIgs Hardware Reference
- o Apple IIgs Technical Note #53, Desk Accessories and Tools
- o Apple IIgs Technical Note #57, The Memory Manager and Interrupts
- o Apple IIgs Technical Note #69, The Ins and Outs of Slot Arbitration

Apple IIgs
#72: QuickDraw II Quirks

Revised by: Dave Lyons
Written by: Dave Lyons & C. K. Haun <TR>

May 1992
November 1989

This Technical Note points out some things things you need to watch out for when using QuickDraw II, especially with FastPort-aware and Shadowing modes.

CHANGES SINCE NOVEMBER 1990: Removed some obsolete information on ScrollRect and on shadowing. Noted that DrawPicture in 6.0 is now compatible with FastPort mode. Added a warning about making QuickDraw II calls while QuickDraw II is not started.

DON'T CALL QUICKDRAW II WHILE IT'S INACTIVE

Most QuickDraw II functions behave unpredictably if you call them while QuickDraw II is inactive, so watch it! Don't make QuickDraw II calls while QuickDraw II isn't started, except as documented. GrafOn and GrafOff are okay. (And so are QDStartUp, QDVersion, and QDStatus.)

FASTPORT-AWARE ANOMALY

Before System 6.0, when the FastPort-aware bit is turned on in the MasterSCB parameter to QDStartUp, DrawPicture did not notice changes in the pen pattern. If your application does not require 6.0 and uses pictures, either directly or indirectly (i.e., by printing to the ImageWriter driver), you may need to leave FastPort-aware mode turned off to get the expected behavior.

FASTFONT AND LARGE PIXEL MAPS

FastFont does not work correctly when drawing past the first 64K of a pixel map. If you are drawing text that uses FastFont (i.e., Shaston 8), you can avoid this problem by using a non-rectangular clipRgn.

DON'T SHOWPEN WHILE COLLECTING POLYGONS, REGIONS, OR PICTURES

The Macintosh QuickDraw documentation permits calling ShowPen after an OpenPoly, OpenRgn, or OpenPicture call to cause drawing calls to contribute to a polygon, region, or picture AND draw to a pixel map at the same time.

The Apple IIgs QuickDraw II documentation does not say you can do that. In some cases, it works, but it works "by accident" and it's not one of the things Apple tests or guarantees in QuickDraw II.

YOU MAY NEED SETBUFDIMS!

The call description for SetBufDims on page 16-215 of Volume 2 of theToolbox Reference is misleading. The note in the description states, "You only need to make this call if your application is going to use, or allow the user to choose, fonts that have unusually large values of chExtra and spExtra." This is not true; you need to call SetBufDims to adjust the clipping buffers for

your application if you plan to use a clipRgn that has a greater width than the width you passed at QDStartUp.

SetBufDims sets the clipping buffer width as well as that of the text buffer, so if you plan to use a clipping region larger than the startup port width you must use SetBufDims.

Be aware that this call may be necessary even if your application does not ever set a clipping region or rectangle. Some toolbox calls assume that the clipping buffer size is correct based on the parameters passed to that routine. For example, if the locInfo you pass to CopyPixels has a width parameter that is wider than the width you passed at QDStartUp, CopyPixels may fail. A safe rule of thumb is to make sure (possibly by setting) that the width parameter in the buffer dimensions is the same or greater than the widest width in the locInfo structures passed to routines that use them.

Further Reference

- o Apple IIgs Toolbox Reference, Volumes 1 and 3

Apple IIgs
#73: Using User Tool Sets

Revised by: Dave "Flag Bits" Lyons
Written by: Dave Lyons

July 1991
November 1989

This Technical Note explains how to write a user tool set and why writing a user tool set is better than stealing a system tool set number.

Changes since January 1991: Expanded recommendation on where to keep user tool set files on disk and clarified SetTSPtr information.

The Apple IIgs Toolbox Reference describes system tool sets, which are usually called through the system tool dispatcher vectors 1 (\$E10000) and 2 (\$E10004).

There are 255 possible system tool set numbers (1 through 255). All of these are reserved for definition by the system. If your program is "borrowing" a system tool set number, please feel guilty and switch over to the user tool set numbers. There are 255 of them too, and they're called through user tool dispatcher vectors 1 (\$E10008) and 2 (\$E1000C). All 255 user tool set numbers are available for the current application to use as it chooses. (Desk accessories are forbidden to use user tool sets.)

Of the four tool dispatcher vectors, only the first one (\$E10000) has received a lot of publicity. \$E10008 works just like \$E10000, except that it passes control to a user tool set instead of a system tool set.

The second vector of each pair (\$E10004 and \$E1000C) works just like the first, except that one extra RTL address must be pushed onto the stack after any parameters are pushed. This way you can have a subroutine to do some or all of your toolbox dispatching, and that subroutine can do extra processing before or after the tool call, or both.

How Can I Write a User Tool Set?

Appendix A of Toolbox Reference, Volume 2, shows how to write a user tool set. Your tool set's Work Area Pointer is a four-byte value you can set with SetWAP and get with GetWAP. The WAP value is already loaded into the Y and A registers every time one of your tool set's functions gets control. The traditional use for the WAP is to keep track of an area of memory owned by your tool set.

If you do use the WAP in a conventional way, your xxxStatus function should return TRUE if the WAP is nonzero; your xxxStartUp function should set the WAP to a non-zero value pointing to some memory space you own (provided by the caller, or allocated with NewHandle using a memory ID provided by the caller); and your xxxShutDown function should set the WAP back to zero.

Since the X register contains the tool set and function number when one of your functions gets control, it is not necessary for a tool set to be written to be used as a predetermined user tool set number. At execution time, your tool set can compute the proper error codes and values to send to GetWAP and

SetWAP.

Note: At the bottom of page A-8 of the Apple IIgs Toolbox Reference, Volume 2, "lda #\$90" should read "lda #\$8100" for version 1.0 prototype. On page A-10, the figure should show two RTL addresses (6 bytes) on the stack.

ToStrip and ToBusyStrip Vectors

These two vectors are for tool sets to jump to when a function exits.

ToBusyStrip \$E10180
ToStrip \$E10184

Inputs: X = error code (0 if no error)
 Y = number of bytes of input parameters to strip

When your function is ready to exit, set up the registers and jump to ToStrip. It shifts the six bytes of RTL addresses up by Y bytes, sets up A and the carry flag appropriately, and returns to whomever called the tool.

If the system busy flag needs to be decremented, jump to ToBusyStrip instead of ToStrip.

How Can I Load My Tool Set From Disk?

One way to load your tool set from disk is to use InitialLoad or InitialLoad2, supplying a pathname like "9:MyToolset" (prefix 9 is initially set to the directory containing your application; prefix 1 also works, but its length is limited to 64 characters). You can then use SetTSPtr to tell the Tool Locator about your tool set, as shown in Appendix A.

Note that SetTSPtr calls your xxxBootInit function. Even if there is no useful work to be done at BootInit time, you still need to have a BootInit function (function number 1) that returns \$0000 in the Accumulator and the carry flag cleared..

When you're done with your tool set, call UserShutdown on the memory ID returned by InitialLoad, so the memory it's using is disposed of or made purgeable. (You can shut it down and allow it to remain in memory in a purgeable state; if you do this, you should try to revive your tool set with Restart before you try InitialLoad or InitialLoad2.)

To allow several applications to share one copy of a user tool set file, you may want to keep your user tool set in the user's *:System:Tools folder. To avoid duplicate file names, leave the ToolXXX names for System tool sets, and give your user tool set a descriptive name.

If your tool set is not found in the *:System:Tools folder, you can then check the 9: folder. This way users do not need to burden their *:System:Tools folders if few of their applications use a particular user tool set or if space on their boot volume is limited.

When your application quits and calls TLShutDown, the system disconnects your tool set from the user tool set TPT. If the UserShutDown is not followed immediately by the TLShutDown, you may wish to use SetTSPtr to cleanly remove your tool set from the system (set the tool set pointer so that it points at a zero word).

Note: Because of the way the tool dispatcher transfers control to toolbox functions, a function's entry point must not be at the first byte of a bank

(\$xx0000). This is normally not an issue, since it's common to put the actual code right after the function pointer table, all in one load segment. Just make sure no function begins at the first byte of a load segment, and you're safe.

Further Reference

- o Apple IIgs Toolbox Reference, Volume 2
- o GS/OS Reference

Apple IIgs
#74: Top Ten List Manager Things

Revised by: Dave Lyons
Written by: Jim Mensch

May 1992
November 1989

This Technical Note presents a method for speeding up custom List Draw routines, with sample source code for the APW assembler.

CHANGES SINCE NOVEMBER 1989: Added information on memFlag and on shared rListRef resources, and noted that System 6.0 already checks the clip region and calls your listDraw routine only when needed.

TEN--MORE MEMFLAG BITS

In each member record, bits 0 and 1 of memFlag indicate whether memPtr is a pointer, handle, or resource ID. You don't normally have to worry about that--a custom listDraw routine is one place that you do. The complete definition of memFlag is as follows:

Bit	Description
7	memSelected
6	memDisabled
5	memNever (Inactive)
4-2	reserved--set to zero
1-0	00 = memPtr is a pointer
	01 = memPtr is a handle
	10 = memPtr is a resource ID (type is rPString or rCString)
	11 = reserved

NINE--SHARING RLISTREF RESOURCES

When listRef is a resource ID, the List Manager calls LoadResource every time it needs your rListRef resource. If two or more lists share the same rListRef, they will get the same handle from LoadResource and will interfere with each other.

To give each list its own copy of your the rListRef resource, load the resource yourself and use DetachResource. Then feed the listRef to the List Manager as a handle. Repeat the process for each list.

EIGHT--CUSTOM LISTDRAW ROUTINES AND THE CLIP REGION

The custom listDraw routine below speeds up your list when running System Software earlier than 6.0. The System 6.0 List Manager already calls your listDraw routine only for members that will not be completely clipped (but this is still a good starting point if you're writing a custom listDraw routine for some other reason).

To scroll text, the List Manager calls ScrollRect to scroll the list--then 6.0 redraws the newly-exposed members, and older versions redraw all the visible members. On small lists this is fine, but on larger lists it can cause the redrawing of much data that is already on the screen, which can take time. If

your application does not require 6.0, you may want to use a custom listDraw routine like this one.

First, we check the current clipRgn (which the List Manager was kind enough to shrink down to include only the portion of the list that needs redrawing) against the passed item rectangle. If the rectangle is in any way enclosed in the clipRgn, then the member is redrawn; otherwise the routine simply returns to the List Manager without drawing. This sample routine is designed to work only with Pascal-style strings, but it can be easily modified to use any other type of string you choose.

```
MyListDraw      Start
;
; This routine draws a list member if any part of the member's
; rectangle is inside the current clipRgn.
;
; Note that the Data Bank register is not defined on entry
; to this routine. If you use any absolute addressing, you
; must set B yourself and restore its value before exiting.
;
top              equ 0
left            equ top+2
bottom         equ left+2
right          equ bottom+2
rgnBounds      equ 2
;
oldDPage       equ 1
theRTL         equ oldDPage+2
listHand       equ theRTL+3
memPtr         equ listHand+4
theRect        equ memPtr+4
using globals

                phd
                tsc
                tcd

                pha
                pha
                _GetClipHandle
                PullLong listHand

                ldy #2
                lda [listhand],y
                tax
                lda [listhand]
                sta listhand
                stx listhand+2

                lda [therect]                ; now test the top
                dec a                          ; adjust and give a little slack
                ldy #rgnbounds+bottom
                cmp [listhand],y             ; rgnRectBottom>=top?
                blt skip2
                brl NoDraw                    ; if not don't draw..
Skip2           ldy #bottom                  ; now see if the bottom is higher
than the top

                inc a                          ; give a little slack
                lda [therect],y
                ldy #rgnBounds+top
                cmp [listhand],y
                blt NoDraw
```

```

NoTest      ANOP

             PushLong theRect
             _EraseRect                ; erase the old rectangle

             ldy #left
             lda [theRect],y
             tax
             ldy #bottom
             lda [theRect],y
             dec a
             phx
             pha
             _MoveTo
             ldy #2
             lda [memptr],y
             pha
             lda [memptr]
             pha
             _DrawString

             ldy #4
             lda [memPtr],y
             and #$00C0                ; strip to the 6 and 7 bits
             beq memDrawn              ; if they are both 0 the member is drawn
             cmp #$0080                ; member selected?
             bne noSelect              ; member not selectable
             PushLong theRect
             _InvertRect
             bra memDrawn

; if we get here the member is disabled
noSelect    PushLong #DimMask
             _SetPenMask
             PushLong theRect
             _EraseRect
             PushLong #NorMask
             _SetPenMask

memDrawn    ANOP

; exit here
           pld
           sep #$20
           longa off
           pla
           ply

           plx
           plx
           plx
           plx
           plx
           plx
           phy
           pha
           rep #$20
           longa on
           rtl

DimMask     dc  i1'$55,$AA,$55,$AA,$55,$AA,$55,$AA'
NorMask     dc  i1'$FF,$FF,$FF,$FF,$FF,$FF,$FF,$FF'
end

```

SEVEN THROUGH ONE--RESERVED FOR FUTURE EXPANSION

Further Reference

- o Apple IIgs Toolbox Reference, Volumes 1 and 3

Apple IIgs
#75: BeginUpdate Anomaly

Revised by: Dave Lyons
Written by: Eric Soldan

May 1992
January 1990

This Technical Note discusses a Window Manager anomaly with the handling of the visRgn and the updateRgn between BeginUpdate and EndUpdate calls.

CHANGES SINCE JANUARY 1990: Updated for System 6.0. CopyPixels is in a static segment, and GS/OS automatically prompts for disks on the text screen when necessary to avoid interfering with a window update in progress.

If an application calls BeginUpdate, it needs to be fully aware of what is going on behind the scenes in terms of its visRgn and updateRgn. Typically an application has TaskMaster handle the update events. TaskMaster calls BeginUpdate, the application update procedure, then EndUpdate. So any application that uses TaskMaster to handle updates, whether or not it makes any BeginUpdate calls directly, needs to be aware of problem described in this Note.

BeginUpdate is responsible for intersecting the visRgn and the updateRgn and making the intersection of these two regions the temporary visRgn. (EndUpdate undoes this effect.) Following are the steps BeginUpdate takes to do this:

1. Localize the updateRgn. (All grafPort regions are local, therefore the visRgn is local. All window regions are global, therefore the updateRgn is global. One of them has to change if they are to be intersected correctly.)
2. Intersect the visRgn and localized updateRgn, then place the result in the updateRgn.
3. Swap the visRgn and updateRgn handles.

The handle swapping has two effects:

- o Makes the intersection region the current visRgn.
- o Saves the real visRgn as the updateRgn. (Saving the real visRgn is necessary because everything has to be restored to normal by EndUpdate.)

EndUpdate restores things to normal after an update procedure is finished. When an application calls EndUpdate, it swaps back the handles and sets the updateRgn to empty.

SO WHAT'S THE PROBLEM?

The problem is that the updateRgn is not a very good place to save the visRgn. Since InvalRect and InvalRgn modify the updateRgn, if either of these two calls is made between a BeginUpdate and EndUpdate, they modify the saved visRgn. When the update is finished, EndUpdate restores the modified visRgn instead of the original.

The solution to this problem seems simple enough: don't call InvalRect or InvalRgn between BeginUpdate and EndUpdate. Unfortunately, there are other calls which can call BeginUpdate, EndUpdate, InvalRect, and InvalRgn, so an

application might inadvertently call one of these routines.

If this situation isn't bad enough already, you could really mess things up by opening another window between BeginUpdate and EndUpdate calls. Opening a window at this time may seem like a perfectly normal thing (i.e., to display an alert); however, opening a window forces the recalculation of the visRgn for any windows obscured by the new window. If the window being updated has its visRgn recalculated, the application obviously loses the visRgn that BeginUpdate created. This doesn't seem too serious since the visRgn is restored to the entire visible part of the window when the new window is closed; however, it does mean that the application would have to update the entire window instead of the original updateRgn.

Unfortunately, the Window Manager also posts update events for the portion of the window that was obscured, and it does this by changing the updateRgn. Of course the updateRgn for the window being updated is really the visRgn that is being "safely" preserved until the EndUpdate call. So, there are some really good reasons why this can't be done.

Okay, so along with not making calls to InvalRect and InvalRgn between BeginUpdate and EndUpdate, an application cannot open any other windows either. Good.

NOW TO MAKE THINGS EVEN WORSE.

Starting with System 5.0, some toolbox functions are stored on disk in dynamic segments and loaded when they are first called. For example, CopyPixels is in a dynamic segment in System versions 5.0 through 5.0.3. If the startup disk is not available and the system prompts for it between BeginUpdate and EndUpdate by calling AlertWindow, the bad things discussed above happen.

Starting with System 6.0, the system is smart enough not to prompt for a disk using AlertWindow if a window update is in progress. (Internally, GS/OS calls WindStatus to see if it can prompt on the graphics screen. If BeginUpdate has been called more times than EndUpdate, WindStatus fibs by returning with the carry set. GS/OS takes the hint and prompts for the disk with a text dialog instead.)

But I Have to Do...

If you absolutely must do some of the things previously discussed, there is a way to accomplish it. It is not simple, but it can be done.

Assuming that BeginUpdate has been called, and an application is in its update procedure:

1. Create a new region and copy the visRgn into it. Doing this allows the application to restore the visRgn to just the area to be updated that BeginUpdate calculated. This needs to be done for any other windows which obscure a part the the window being updated. Again, these are not windows that an application would open directly. CopyPixels may open a window, since it is a dynamic segment and may need to get loaded from a disk that is off-line.
2. Create a new region, then swap its handle with the updateRgn handle. This protects the real visRgn and lets an application call InvalRect and InvalRgn at any time if necessary. It also means the application doesn't need to worry about the system making these calls either. The updateRgn is also an empty region after the swap, so any contributions to it constitute a valid update event that needs to be handled.
3. Do the update part of the update procedure. In this part, if the application has any calls to CopyPixels, or any other QuickDraw

Auxiliary dynamic segment functions, after the call is completed, copy the saved visRgn back to the visRgn of the grafPort. The closing of the dynamic segment alert window recalculates the visRgn, and copying it undoes this effect. Do not do the same for the updateRgn. Leave the updateRgn alone. We are accumulating an actual updateRgn, and the closing of the alert window for the dynamic segment may have contributed to this region.

There are two methods for leaving the update procedure. Although the second method works whether or not an application uses TaskMaster, if an application does not use TaskMaster, then the first method is simpler.

The procedure without using TaskMaster (i.e., you made the BeginUpdate call, and you will make the EndUpdate call) is as follows:

- A. Dispose of the region created in Step 1. This region was only needed to restore the partial visRgn that BeginUpdate calculated after a window was opened.
- B. Swap the updateRgn handle with the region handle created in Step 2.
- C. Make the EndUpdate call.
- D. If the region created in Step 2 is not empty, copy this region into the updateRgn for the window with CopyRgn. You can't just do an InvalRgn with it because InvalRgn globalizes the region and the region is already global. You want to copy the region since this generates a valid update event. You can use CopyRgn instead of UnionRgn because the update region is empty.
- E. Dispose of the region created in Step 2.

With TaskMaster, things are a little messier. Since TaskMaster makes the EndUpdate call, you have less control over the situation. It is important to do the EndUpdate before generating the update event. Posting the update event has to happen outside the update procedure, since you have to leave the update procedure for TaskMaster to do the EndUpdate. So it follows that you do Steps A and B, post an application event to handle the rest externally, and when the application event is handled, do Steps D and E.

Some consideration was given to posting an application event via the PostEvent call. Unfortunately, there is a possibility that this application event will drop out of the queue not handled. When the queue is full, the oldest event is dropped, and this could occur to application events, which would be very bad in this case. Due to this possibility, posting an application event refers to setting a global variable that is checked before the TaskMaster call in the main event loop. This can be considered equivalent to posting an event via the PostEvent call.

So, the TaskMaster case would be as follows:

- A. Dispose of the region created in Step 1.
- B. Swap the updateRgn handle with the region handle created in Step 2.
- C. Store the handle of the region created in Step 2 in a global variable named eventUpdateRgn. Store the current window port in a global variable named eventWindowPort.
- D. Return to TaskMaster, which returns to the main event loop.
- E. Immediately after the TaskMaster call in the main event loop, check the global variable eventUpdateRgn. If it is not NULL then:
 - a. Copy the region into the updateRgn of the window eventWindowPort. Using CopyRgn is the easiest way to copy the region. (Copying the region posts an update event if the eventUpdateRgn is not NULL.

- b. Dispose of the region eventUpdateRgn, then set the variable eventUpdateRgn to NULL, so that this "event" won't be handled again.

Of course, the simplest way to handle all of this is to avoid situations where you have to take the steps described above. If things like opening a window (or allowing the system to open one) and InvalRect and InvalRgn can be avoided between calls to BeginUpdate and EndUpdate, so can all of this ugliness.

Further Reference

- o Apple IIgs Toolbox Reference, Volume 2

Apple IIgs
#76: Miscellaneous Resource Formats

Revised by: Matt Deatherage
Written by: Matt Deatherage, C.K. Haun, Llew Roberts
& Dave Lyons

May 1992
January 1990

This Technical Note describes resource structure formats for previously-unpublished types.

CHANGES SINCE DECEMBER 1991: Added information on rFont resources. Clarified the note about rVersion resources to note that version numbers must increase with subsequent releases for the Finder's benefit.

The format used to describe the resources is similar to that used in File Type Notes, where the offsets, given in the form (+xxx), determine the offset from the beginning of the resource.

SAMPLED SOUND RESOURCE (TYPE: \$8024, RSOUNDSAMPLE)

The following describes the Sampled Sound resource format. It consists of a ten-byte header followed by the sample data bytes.

Format	(+000)	Word	This must always be zero.
Wave Size	(+002)	Word	Sample size in pages (256 bytes per page). For example, an 8K sample takes 32 pages; a 128K sample requires \$200 pages.
Rel Pitch	(+004)	Word	The high byte of this word is a semitone value; the low byte is a fractional semitone. These values are used to tune the sample to correct pitch. See HyperCard IIgs Technical Note #3, Pitching Sampled Sound.
Stereo	(+006)	Word	The output channel for this sound is in the low nibble of this word.
Sample rate	(+008)	Word	The sampling rate of the sound, in Hertz (Hz).
Sound	(+010)	Bytes	The sampled sound data. The bytes are all 8-bit samples. The sample starts here and continues until the end of the resource.

The resource compiler template follows:

```
#define rSoundSample      $8024

/*----- rSoundSample -----*/
type rSoundSample {
    integer;           /* format */
    integer;           /* wave size */
    hex integer;      /* rel pitch */
    integer;           /* stereo channel */
    unsigned integer; /* sample rate */
    hex string;       /* raw 8 bit sound data */
};
```

CURSOR RESOURCE (TYPE: \$8027, RCURSOR)

The following describes the Cursor resource format:

height (+000) Word The height of the cursor, in pixels.
width (+002) Word The width of the cursor, in Words.
image (+004) Bytes The image of the cursor. There are height*width
Words in the cursor, or twice that many Bytes.

We will call the first byte beyond the image offset "ei" for "end of image."

mask (+ei) Bytes The mask of the cursor. This is the same size as the image.

We will call the first byte beyond the mask offset "em" for "end of mask."

hotSpotY (+em) Word The cursor's Y "hot spot."
hotSpotX (+em+2) Word The cursor's X "hot spot."
flags (+em+4) Flag Cursor flags:
Bit 7: 1 = 640 Mode, 0 = 320 Mode
All other bits are reserved and must be zero.
reserved (+em+6) 8 Bytes Reserved, must be zero.

The resource compiler template follows:

```
#define rCursor      $8027

/*----- rCursor -----*/
type rCursor {
    height :
        hex integer;          /* height */
    width :
        hex integer;          /* width in words */
        hex string[2*$$Word(height)*$$Word(width)]; /* cursor image */
        hex string[2*$$Word(height)*$$Word(width)]; /* cursor mask */
        hex integer;          /* hotspot Y */
        hex integer;          /* hotspot X */
        hex integer;          /* flags */
        hex longint = 0;      /* reserved */
        hex longint = 0;      /* reserved */
};
};
```

Following is a simple cursor example:

```
resource rCursor(1, fixed) {
    5, /* height */
    2, /* width */
    $"ffff0000"
    $"f00f0000"
    $"f00f0000"
    $"f00f0000"
    $"f00f0000"
    $"ffff0000",

    $"ffff0000"
    $"ffff0000"
    $"ffff0000"
    $"ffff0000"
    $"ffff0000",

    2, /* hotspot Y */
    2, /* hotspot X */
};
```

```
    $80    /* 640 mode */  
};
```

Note that the resource is marked fixed so that its handle can be dereferenced and passed to SetCursor.

VERSION RESOURCE (TYPE: \$8029, RVERSION)

Files may include a version resource with ID=1 for display by programs such as the Finder. All rVersion resource IDs other than 1 are reserved for future definition. The following describes the version resource format:

version	(+000)	Long	The application's version number, in Apple IIgs Long Version format. See Apple IIgs Technical Note #100, VersionVille, for more details.
country	(+004)	Word	An international country version code. Possible values are as follows:

verUS	0
verFrance	1
verBritain	2
verGermany	3
verItaly	4
verNetherlands	5
verBelgiumLux	6
verSweden	7
verSpain	8
verDenmark	9
verPortugal	10
verFrCanada	11
verNorway	12
verIsrael	13
verJapan	14
verAustralia	15
verArabia	16
verFinland	17
verFrSwiss	18
verGrSwiss	19
verGreece	20
verIceland	21
verMalta	22
verCyprus	23
verTurkey	24
verYugoslavia	25
verIreland	50
verKorea	51
verChina	52
verTaiwan	53
verThailand	54

name	(+006)	String	Pascal string containing the desired name. May be the null string.
moreInfo	(+xxx)	String	Additional information to be displayed, such as a copyright notice. May be the null string. Recommended maximum length is about two lines of 35 characters each. May contain a carriage return (character \$0D).

The resource compiler template follows:

program. No length limit is imposed by this resource format, although a practical limit of a few hundred characters is recommended.

The resource compiler template follows:

```
#define rComment      $802A

/*----- rComment -----*/
type rComment {
    string;
};
```

TAGGED STRINGS RESOURCE (TYPE: \$802E, RTAGGEDSTRINGS)

A tagged strings resource lists pairs of Word values and Pascal strings.

count	(+000)	Word	Number of word/string pairs in this resource.
firstWord	(+002)	Word	Word value of first pair.
firstString	(+004)	String	Pascal string of first pair.
secondWord	(+xxx)	Word	Word value of second pair.
secondString	(+yyy)	String	Pascal string of second pair.
...			

The resource compiler template follows:

```
#define rTaggedStrings      $802E

/*----- rTaggedStrings -----*/

type rTaggedStrings {
    integer = $$Countof(StringArray);
    array StringArray {
        hex integer;          /* Key integer */
        pstring;             /* String */
    };
};
```

Following is a simple rTaggedStrings example:

```
resource rTaggedStrings(1) {{
    $0050, "red",
    $0033, "green",
    $0100, "blue"
}};
```

PATTERN LIST RESOURCE (TYPE: \$802F, RPATTERNLIST)

A pattern list resource contains zero or more 32-byte QuickDraw II patterns. (This resource type exists for your convenience. The System Software contains no direct support for resources of this type.)

firstPattern	(+000)	32 Bytes	First QuickDraw II pattern structure.
secondPattern	(+032)	32 Bytes	Second QuickDraw II pattern structure.
...			

The resource compiler template follows:

```

#define rPatternList          $802F

/*----- rPatternList -----*/

type rPatternList {
    array {
        array[32] {
            hex byte;
        };
    };
};

```

RECTANGLE LIST RESOURCE (TYPE: \$C001, RRECTLIST)

The rectangle list (type rRectList) is provided to allow an extensible, easily modifiable collection of QuickDraw II rectangle structures. This capability can enhance a developer's ability to modify on-screen displays without recompiling an entire application. This resource also enables easier cross-development and parallel development for the Apple IIgs and the Macintosh.

The following describes the rectangle list resource format:

count	(+000)	Word	Number of rectangles in this resource.
firstRectangle	(+002)	8 Bytes	First QuickDraw II rectangle structure.
	...		Rectangles, eight bytes each.
lastRectangle	(+002+(8*(count-1)))	8 Bytes	Last QuickDraw II rectangle structure.

The resource compiler template follows:

```

#define rRectList          $C001
type rRectList {
    integer = $$Countof(RectArray);
    array RectArray {
        Rect;
    };
};

```

PRINT RECORD RESOURCE (TYPE: \$C002, RPRINTRECORD)

As a convenience for applications, a print record may be included as a resource of type \$C002 (rPrintRecord). If more than one of these resources is present, the one to use as the document's primary print record is the first one. You can get this resource's ID by calling GetIndResource with type rPrintRecord and index 1. Storing the primary print record with ID = 1 is a good way to start.

Since the print record is filled in and interpreted by the printer driver, you can't always programmatically set options that are driver-specific. For example, although the ImageWriter driver stores the color-vs.-black and white option in one place, not all color printers will do the same thing. If you want to use driver-specific options on many printers, you can use those printer drivers to create print record that reflect the options you want and store those records as resources. Then, if you need a pre-initialized print record with the options you want, you may already have one.

print record	(+000)	160 Bytes	The print record. The handle to this resource is suitable for passing to any
--------------	--------	-----------	--

Print Manager call that requires a print record handle.

Since applications shouldn't create print records from scratch, but rather allow printer drivers to fill them in with PrDefault and PrVerify, the resource compiler template that follows only allocates 160 bytes for storage.

```
#define rPrintRecord    $C002

/*----- rPrintRecord -----*/
type rPrintRecord {
    array[160] {
        hex byte;
    };
};
```

FONT RESOURCE (TYPE: \$C003, RFONT)

Some applications wish to keep fonts with the application itself instead of in a separate font file. This isn't always advisable--fonts not in font files can't be easily used in other applications, which may confuse users who, for example, use text-editing desk accessories and can't get at certain fonts except in certain applications. Also, fonts not in the Fonts directory must be completely memory-resident where normal Fonts are only loaded from disk when they're needed.

Nevertheless, in some cases keeping fonts inside an application file is necessary or desirable. In such cases, the rFont resource is a convenient way to keep a font around. The resource is a QuickDraw II Font, as defined in Apple IIgs Toolbox Reference, Volume 2. The font family name is the resource's name--the same Pascal string that normally precedes the QuickDraw II font record in the font file.

font (+000) Bytes The font record, without the font family name.

Since the font record is a bunch of variable-sized tables, and since you probably want to use a font editor (and not the resource compiler) to create fonts, the following resource compiler template isn't very revealing.

```
/*----- rFont -----*/
type rFont {
    hex string;
};
```

Further Reference

- o Apple IIgs Toolbox Reference, Volumes 1-3
- o Apple IIgs Technical Note #100, VersionVille
- o HyperCard IIgs Technical Note #3, Pitching Sampled Sounds

Apple IIGS

#77: Print Manager & AppleTalk Configuration Files

Written by: Jim Luther

January 1990

This Technical Note describes the Print Manager user configuration file `Printer.Setup` and the AppleTalk user configuration file `ATInit`. This Note also describes a limitation of the Print Manager call `PrGetUserName`, which is a result of the way configuration information is stored in the `Printer.Setup` file.

Printer.Setup and ATInit

What Are the Printer.Setup and ATInit Files?

The Print Manager user configuration file `Printer.Setup`, which is found in the `System:Drivers` directory of the Apple IIGS boot disk, is used by the Print Manager tool set to keep the name of the printer driver and port driver you've selected between system boots. In addition, if you've selected a network printer, `Printer.Setup` contains the selected printer's network address and your machine's User Name. The file format of `Printer.Setup` has not been published because revisions have been made, and may be made again, to the Apple IIGS System Software, which can change `Printer.Setup`'s file format.

The AppleTalk user configuration file `ATInit`, which is found in the `System:System.Setup` directory of the Apple IIGS boot disk, is used to keep the default AppleShare startup application, the default AppleShare prefix, the default AppleTalk User Name, and the default AppleTalk printer entity name (the network printer entity used by AppleTalk's Remote Print Manager) between system boots. The file format of the `ATInit` file was published incorrectly in the AppleShare Programmer's Guide for the Apple IIGS. The correct file format for `ATInit` will be discussed later in this Note.

It is important to remember that the Print Manager tool set uses the information from the `Printer.Setup` file only, and that AppleTalk and AppleShare use the information contained in the `ATInit` file only. It is also important to note that the Print Manger tool set, which is used to print QuickDraw II graphics, and AppleTalk's Remote Print Manager (RPM), which is used to print ASCII data to network printers, are not the same thing even though both contain the words "Print Manager."

What Writes to the Printer.Setup and ATInit Files?

Before Apple IIGS System Software 5.0, `Printer.Setup` and `ATInit` were handled as completely separate configuration files. The Print Manager call `PrChoosePrinter` allowed you to select the printer and port drivers the Print Manager would use and wrote the printer and port driver selections to the `Printer.Setup` file. The AppleTalk application `Chooser.II` let you select the printer AppleTalk's Remote Print Manager would use and wrote the printer entity selection to the `ATInit` file.

With System Software 5.0 all printer selections for both the Print Manager and AppleTalk are made by using one of the Control Panel NDA's printer CDevs. All printer CDevs (e.g., `DirectConnect`, `ATIWriter`, `ATLQIWriter`, and `ATLWriter`)

write the new printer and port driver selections to the Printer.Setup file. However, if the printer selected uses the AppleTalk port (i.e., the selection is made with the ATIWriter, ATLQIWriter, or ATLWriter printer CDevs), then the selected printer's network address and your User Name are written to both the Printer.Setup and the ATInit files. The DirectConnect CDev does not write any information to the ATInit file. If AppleShare is installed, then the AppleShare CDev will also write your User Name to the ATInit file.

On AppleShare file servers with the Apple II Setup option installed, the ATInit file in User folders will also be written to by the AppleShare Admin application when the Apple II startup information is set.

When are the Printer.Setup and ATInit Files Read?

The Printer.Setup file is read by the Print Manager and by the printer CDevs. The Print Manager reads the information contained in the Printer.Setup file whenever the Print Manager needs to load a printer driver or a port driver into memory. A printer CDev reads the information contained in the Printer.Setup file when that CDev is selected so it can know the current printer and port selections.

Ways the printer driver and the port driver might be unloaded and need to be loaded (which will cause Printer.Setup to be read by the Print Manager) are as follows:

- o The Print Manager is shut down.
- o The current printer driver or port driver is changed with a Control Panel printer CDev. When a new printer or port is selected with a printer CDev, the current drivers are unloaded from memory so the Print Manager will be forced to read the new printer and port selections from Printer.Setup.
- o Your application makes the PMUnloadDriver Print Manager call.

An application can load one or both of the drivers (which will cause Printer.Setup to be read by the Print Manager) by making the PMLoadDriver call. The AppleTalk user configuration information contained in the ATInit file is read during system startup as part of AppleTalk's initialization.

Network Booting and Printer.Setup

When Apple IIGS computers are booted over an AppleShare network, they all share a single copy of the Printer.Setup file. That means all machines must use the same printer and port driver selections that are stored in the Printer.Setup file. If all machines are expected to be able to print using the Print Manager tool set, then the printer and port selection stored in Printer.Setup must be something that all can use. The only two options are:

- o A single shared network printer for all machines (i.e., a LaserWriter, an AppleTalk ImageWriter, or an AppleTalk ImageWriter LQ). In situations where many machines are booted over a single file server, this may cause the workload on the single shared printer to be unacceptable.
- o A direct-connect printer on each machine. The limitations of this solution are that the printers must be of the same type (all ImageWriters, all ImageWriter LQs, or all Epsoms) and all machines must use the same printer port (either printer or modem).

The server administrator should set the default printer selection, which will

be used by all machines, by using one of the Control Panel NDA's printer CDevs. Then, the access privileges to the server's System:Drivers directory should be set to "Bulletin Board" (i.e., Everyone See Folders, Everyone See Files, Owner Make Changes) so other machines cannot change the printer and port selection.

Using User Names

The User Name We Use

You may have noticed that you see your AppleTalk User Name in the Control Panel's AppleShare and printer CDevs. AppleShare allows a machine's User Name to be up to 31 characters long. The CDevs read the User Name from the ATInit file. The AppleShare and printer CDevs also store the complete User Name back into the ATInit file.

PrGetUserName (Almost)

The Printer.Setup file sets aside 15 characters for the User Name so the printer CDevs store only the first 15 characters of the User Name in the Printer.Setup file. This limitation is leftover from early Print Manager implementations of the PrChoosePrinter call, which limited the User Name length to 15 characters.

Since the Print Manager gets the User Name it uses from the Printer.Setup file, the User Name returned by the Print Manager call PrGetUserName will be truncated to 15 characters if the complete AppleTalk User Name is 16 characters or longer.

Where to Get the Complete User Name

If your application needs the complete default AppleTalk User Name, it can be read from the ATInit file. When an Apple IIGS is booted from a local disk volume that has AppleShare or at least one of the AppleTalk network printers installed, ATInit will be found in the System:System.Setup directory of the local boot volume. When an Apple IIGS is booted over AppleTalk, ATInit will be found in the Users:YourName:Setup directory of the AppleShare boot volume (where YourName is the User Name used to log on to the boot server).

The ATInit File Format

The AppleShare Programmer's Guide for the Apple IIGS shows the file format of the ATInit file as it is stored on an AppleShare boot volume. However, the file format of ATInit is not always as shown in that manual. In all cases, ATInit will contain the three required data fields UserName, PrinterFlags, and PrinterTuple at the end of the file. Before those data fields, ATInit may also contain executable code or additional data fields. Since the three required data fields are directly before ATInit's end-of-file (EOF), you can find them relative to ATInit's EOF using the displacements listed in Table 1.

Displacement to ATInit EOF	Size	Field Name	Description
133	33 Bytes	UserName	A Pascal-type string containing the default User Name. It consists of a length byte followed by up to 31 bytes of ASCII data and a single, unused byte. This field is always 33 bytes long.

100	Byte	PrinterFlags	This is the Flags field used by the Remote Print Manager's default network printer.
99	99 Bytes	PrinterTuple	This field specifies the name of the default network printer used by the Remote Print Manager. The PrinterTuple field is in standard Name Binding Protocol (NBP) format. This field is always 99 bytes long.

Table 1-Offsets of Required Data Fields

If the ATInit file is on an AppleShare server, it will have 6 additional data fields (PathVolID, PathDirID, Path, PrefixVolID, PrefixDirID, and Prefix) directly before the three required data fields. These fields can also be found relative to ATInit's EOF using the displacements listed in Table 2.

Displacement to ATInit EOF	Size	Field Name	Description
275	Word	PathVolID	The Volume ID number of the user's AppleTalk startup application.
273	Long	PathDirID	The Directory ID number of the user's AppleTalk startup application.
269	65 Bytes	Path	The Pathname of the user's AppleTalk startup application.
204	Word	PrefixVolID	The Volume ID number of the user's AppleTalk default prefix.
202	Long	PrefixDirID	The Directory ID number of the user's AppleTalk default prefix.
198	65 Bytes	Prefix	The user's AppleTalk default prefix.

Table 2-Offsets of Optional Data Fields

The displacements in Tables 1 and 2 can be used with the GS/OS SetMark call to move the file mark to the beginning of any of the above fields. The SetMark call's base field should be set to \$0001 so the mark will be set equal to EOF minus the displacement.

Further Reference

- o Apple IIGS Toolbox Reference
- o Inside AppleTalk
- o AppleShare Programmer's Guide for the Apple IIGS

Apple IIgs
#78: Bank Alignment and Memory Management

Revised by: Matt Deatherage
Written by: Matt Deatherage

May 1992
March 1990

This Technical Note discusses the way the Memory Manager deals with requests for memory that is already in use, and why this can be really annoying.

CHANGES SINCE MARCH 1990: Included new information about some smarter algorithms in System Software 6.0 and later which can avoid problems some of the time.

The Memory Manager is a sophisticated software module that provides the framework for the allocation, moving, management, and disposal of blocks of memory; however, it's not magic.

When you ask the Memory Manager for a block of memory and it's not immediately allocatable, the Memory Manager starts through the procedure for purging, compacting, and calling out-of-memory (OOM) queue routines until, at the end of its rope, it finally gives up and returns error \$0201. The exact procedure is repeated below, taken from Volume 3 of the Apple IIgs Toolbox Reference. Note that each successive step is only taken if, after the previous step, the requested memory still isn't available.

1. Calls each OOM queue routine until either all routines have been called or until one OOM queue routine reports that it has freed enough memory to satisfy the request.
2. Compacts memory.
3. Purges all level 3 handles. If this frees enough memory, compaction occurs.
4. Purges all level 2 handles. If this frees enough memory, compaction occurs.
5. Purges all level 1 handles. If this frees enough memory, compaction occurs.
6. Calls each OOM queue routine again until all have been called or until one OOM queue routine reports that it has freed enough memory to satisfy the request.
7. Gives up and returns error \$0201.

This strategy works pretty well--as long as the request is for a block of memory wherever it fits. If someone has asked the Memory Manager for memory at a specific address, things get stickier.

Suppose that you've asked the Memory Manager for a handle starting at the beginning of bank 2, and that something else (i.e., the ProDOS FST) is already using that memory. The Memory Manager notices that the handle isn't immediately available, so it starts going through the listed procedures. Since the handle for the ProDOS FST is neither purgeable nor movable and GS/OS isn't likely to give it up in an OOM queue routine, the request fails and the Memory Manager returns error \$0201.

However, the Memory Manager went through all the steps listed to get to the seventh step, the error. The Memory Manager has no way to know that one of the OOM queue routines isn't going to give up that particular handle and allow the request to be fulfilled. The OOM queue routines cannot know themselves,

since they are only told how much memory is needed, not where it has to be. Therefore, whenever the Memory Manager returns error \$0201, all purgeable handles have been purged.

This is particularly annoying to loaders. OMF supports a "bank-aligned" attribute for load segments, and the loaders ensure that such segments are loaded at the beginning of some bank or another. The Memory Manager does not have a "bank-aligned" attribute for handles, so the loaders have to do these things themselves. They do this by asking for a handle of the appropriate size at the beginning of bank two. If this fails, the loaders try again with bank three, then bank four, and so on through the end of memory.

Since some part of GS/OS is almost always occupying the memory at the beginning of bank two, which is where the loader first attempts to load a bank-aligned segment, the presence of such a segment in a load file virtually guarantees that all purgeable handles are purged when the file is loaded. This kicks out dormant applications and zombie-state tool sets, among other things, requiring they be loaded from disk again when needed.

Starting with System Software 6.0, the Loader attempts an alternate strategy first--it tries to allocate an entire bank of memory that is page aligned and doesn't cross a bank boundary. This block, if available, will by definition be bank-aligned. Since code segments can't be larger than 64K, such a block can always hold a bank-aligned segment. The Loader now tries to allocate such a block and if it finds one, it immediately disposes of it and allocates a block of the right size at the same bank address.

If this strategy succeeds, it's a lot faster than the other method and may avoid purging all of memory. If it fails, though, the Memory Manager still goes through all seven steps before returning error \$0201, so all of memory may still be purged. It's just less likely in System Software 6.0 and later.

It doesn't make sense to bank-align a small segment, and small segments fit better into fragmented memory. If you use large segments anyway, consider the trade-off: bank-aligning a segment may purge memory at load time, but your linker may be able to generate smaller OMF, decreasing disk size and load time.

SUMMARY

The general recommendation against asking for specific blocks of memory is well-known to most developers; the reasons outlined above simply add fuel to the fire against such programming practices. What isn't as widely known is that having a bank-aligned load segment in a load file may cause everything purgeable to be purged, and could also cause OOM queue routines to dispose of handles when there really isn't any kind of memory shortage.

Apple advises developers to carefully consider the advantages and disadvantages of bank-aligned segments before including one in a load file.

Further Reference

- o Apple IIgs Toolbox Reference, Volumes 1 and 3
- o GS/OS Reference

Apple IIGS
 #79: Integer Math Data Types

Revised by: Jim Luther
 Written by: Dan Strnad

May 1990
 March 1990

This Technical Note describes the format of Fixed and Frac data types used by the Integer Math tool set and operations performed on the Integer Math numerical data types.

Revised since March 1990: Fixed original date, bit numbering of diagrams, and a multiplication sign in the equation.

As stated in Volume 1 of the Apple IIgs Technical Reference, the Integer Math tool set provides the following numerical data types:

- Integers
- Longints
- Fixed
- Frac
- Extended

The precise format of the Fixed and Frac data types is not provided in the reference manual, so this Note details these formats.

The format for the Fixed data type is stated in the manual as being a 32-bit signed value with 16 bits of fraction. This means that the low-order 16 bits of the Fixed format data value are considered as a fraction of 2^{16} , which is the binary number represented by a one followed by 16 zeroes (\$10000). In other words, a Fixed value is the same as a long integer value whose binary point has been moved to the left 16 places. In this representation, if the low-order part of the Fixed format data value were \$8000, the fractional value would be equal to $1/2$. A low-order part of \$C000 would represent a fractional part equal to $3/4$. Therefore the highest value that a Fixed can contain is 32,767 and 65,535/65,536; the least value is equal to -32768.

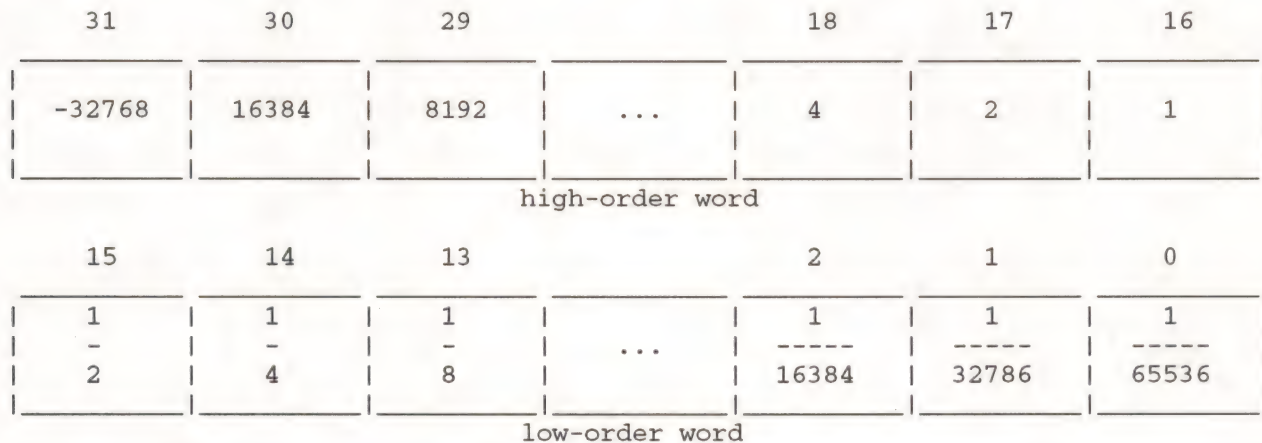


Figure 1-Fixed Data Type

The format for the Frac data type is stated in the manual as being a 32-bit signed value with 30 bits of fraction. This means that the low-order 30 bits of the Frac format data value are considered as a fraction of 2^{30} , which is

the binary number represented by a one followed by 30 zeroes (\$40000000). In other words, a Frac value is the same as a long integer value whose binary point has been moved to the left 30 places. The high-order 2 bits of the Frac format data value are treated as follows. The high bit has a value of -2 and the low bit has a value of 1. Therefore the highest value that a Frac can contain is 1 and $((2^{30})-1)/2^{30}$; the least value is equal to -2.

31	30	29		18	17	16
-2	1	1 - 2	...	1 ---- 4096	1 ---- 8192	1 ----- 16384
high-order word						
15	14	13		2	1	0
1 ----- 32768	1 ----- 65536	1 ----- 131072	...	1 ----- 268435456	1 ----- 536870912	1 ----- 1073741824
low-order word						

Figure 2-Frac Data Type

Note that for Longints, Fixed, and Frac values, the hex representations of the largest and smallest data values are \$7FFFFFFF and \$80000000, respectively.

A property of the Fixed and Frac data types is that two Fixed or two Frac values may be added or subtracted just as if they were 32-bit integers. To demonstrate this, imagine scaling the numbers by a given factor to make them integers. After adding the numbers, the sum could be scaled back down by the same factor. This follows from the distributive property of multiplication over addition, which allows one to make the inference shown in the equations which follow. In these equations, V1 and V2 are both either Fixed or Frac values. The value for C being discussed, which illustrates the ability to scale Fixed and Frac values, is 2^{16} for Fixed values of V1 and V2, or 2^{30} for Frac values of V1 and V2.

$$\frac{(C * V1) + (C * V2)}{C} = \frac{C * (V1 + V2)}{C} = V1 + V2$$

Similarly, two Fixed or two Frac values may be compared, as Longints are compared, with one another. In general, the comparison, addition, and subtraction operations used for long integers may also be performed on any two Fixed or any two Frac values.

Further Reference

-
- o Apple IIGS Technical Reference Manual
 - o Apple Numerics Manual, Second Edition

Apple IIGS
#80: QuickDraw II Clipping

Written by: Eric Soldan March 1990

This Technical Note explains a lot about QuickDraw II operation, specifically clipping.

Before Beginning

Before beginning this Note, some statements, disclaimers, and definitions:

1. This is not a substitute for the QuickDraw II introduction in the Apple IIGS Toolbox Reference, but rather a supplement.
2. A pixelmap is a series of bytes that hold pixel data whose rectangular shape is defined by a LocInfo structure.

This Note describes in great detail the way that QuickDraw II does things with pixelmaps. It begins with a description of the LocInfo structure, which is the most important thing to understand in terms of QuickDraw II pixelmap management. Once this is understood, this Note covers how it applies to using functions such as PPToPort, PaintPixels, and CopyPixels. And once this is understood, it then describes how LocInfo structures are used to control drawing into a grafPort. (PPToPort is used in this Note. PaintPixels and CopyPixels are very close in function to PaintPixels. The information and theory in this Note also apply to these calls.)

Understanding the material in this Note should help you better understand the entire toolbox. It is surprising how much can be accomplished with the toolbox without completely understanding these concepts; it is also surprising how much easier programming with the toolbox gets when these concepts are fully understood.

Note: Structures are written with C syntax in this Note. In addition, this Note uses the screen address 0xE12000L. The possibility of shadowing being active and the screen address being 0x12000L is ignored.

The Beginning

One must begin with the LocInfo structure, which is as follows:

```
struct LocInfo {
    Word      portSCB;           /* SCB in low byte */
    Pointer   ptrToPixImage;    /* ImageRef */
    Word      width;            /* Width */
    Rect      boundsRect;      /* BoundsRect */
};
```

For this Note, one can change this structure a little bit by calling the width element rowBytes. This convention is good because rowBytes is more descriptive than width (it indicates that one is measuring the width in bytes) and it allows one to use the word "width" elsewhere in this Note without confusion. So for the purposes of the Note, the new LocInfo structure definition is as follows:

```

struct LocInfo {
    Word      portSCB;          /* SCB in low byte */
    Pointer   ptrToPixImage;   /* ImageRef */
    Word      rowBytes;        /* Width in bytes*/
    Rect      boundsRect;      /* BoundsRect */
};

```

The ptrToPixImage field is a pointer to some block of bytes in memory. (This block of bytes is referred to as the pixImage from here on.) A pixImage doesn't have any inherent shape. QuickDraw II deals with it as a rectangle, and the LocInfo record defines the rectangularity of it.

When saving a 32,000 byte screen image, one doesn't save the number of bytes of which each row consists. One assumes that each row is 160 bytes by convention, and this is a safe assumption, since the IIGS video hardware expects 160 bytes. But the point is that in the 32,000 bytes of screen data, there is no indicator as to the specific size of a row. One must just know that it is 160 bytes per row. This size is fine for screen shots, but it is not fine when different pixelmaps can be different widths. If they can be different widths, then one also needs some information as to what those widths are, hence the portSCB, rowBytes, and boundsRect fields in a LocInfo structure.

The boundsRect and portSCB fields tell the shape of the pixelmap in pixels, the boundsRect tells how many pixels wide and tall the pixelmap is, and the portSCB tells how big those pixels are (320-mode pixels are four bits wide and 640-mode pixels are two bits wide). One would think that this would be enough information to determine the size of the pixImage, but it isn't. The rowBytes can be larger than the boundsRect/portSCB would indicate (see Figure 1). This situation is legal; it means that some bytes are being wasted, but it is legal.




Figure 1-Sample LocInfo Structure

One simply has to know the size of the pixImage, since it cannot be determined by the LocInfo information. If the pixImage is the screen, then it is 32,000

bytes. If it is a fixed or locked handle, then one can do a FindHandle on the pointer followed by a GetHandleSize on the found handle.

Figure 1 represents a sample LocInfo structure. The portSCB (although not pictured) is also relevant, as it determines the size of the pixels. If the pixelmap is a 320-mode pixelmap, one could change it to a 640-mode pixelmap by changing the portSCB to 640 mode and doubling the width of the boundsRect. In doing this conversion, note that rowBytes is not affected and that the pixImage does not change size.

In the example illustrated in Figure 1, the pixImage is bigger than the boundsRect, but again, this is okay. However, this is not the case for the screen, where the rowBytes is 160 and the height of the boundsRect is 200 (the size of the screen is exactly equal to $160 * 200 = 32,000$).

There are some rules to determining the rowBytes value. First, rowBytes must not be too small. This is obvious. Second, rowBytes must be evenly divisible by eight. This is not at all obvious, but it is very important. QuickDraw II makes some assumptions for speed, and one of them is that rowBytes is a multiple of eight.

So much for describing the LocInfo structure. Now for how to use it via PPToPort.

PPToPort accepts (among other things) a pointer to a source LocInfo record and a pointer to a source rectangle. PPToPort does not use the source rectangle directly; it first intersects it with the boundsRect in the LocInfo record, and it uses this intersection rectangle instead. This intersection rectangle guarantees that the area involved is completely enclosed by the boundsRect (and therefore within the pixImage). If the source rectangle is entirely outside the boundsRect, then the intersection of the source rectangle and the boundsRect is empty, thus nothing is drawn.

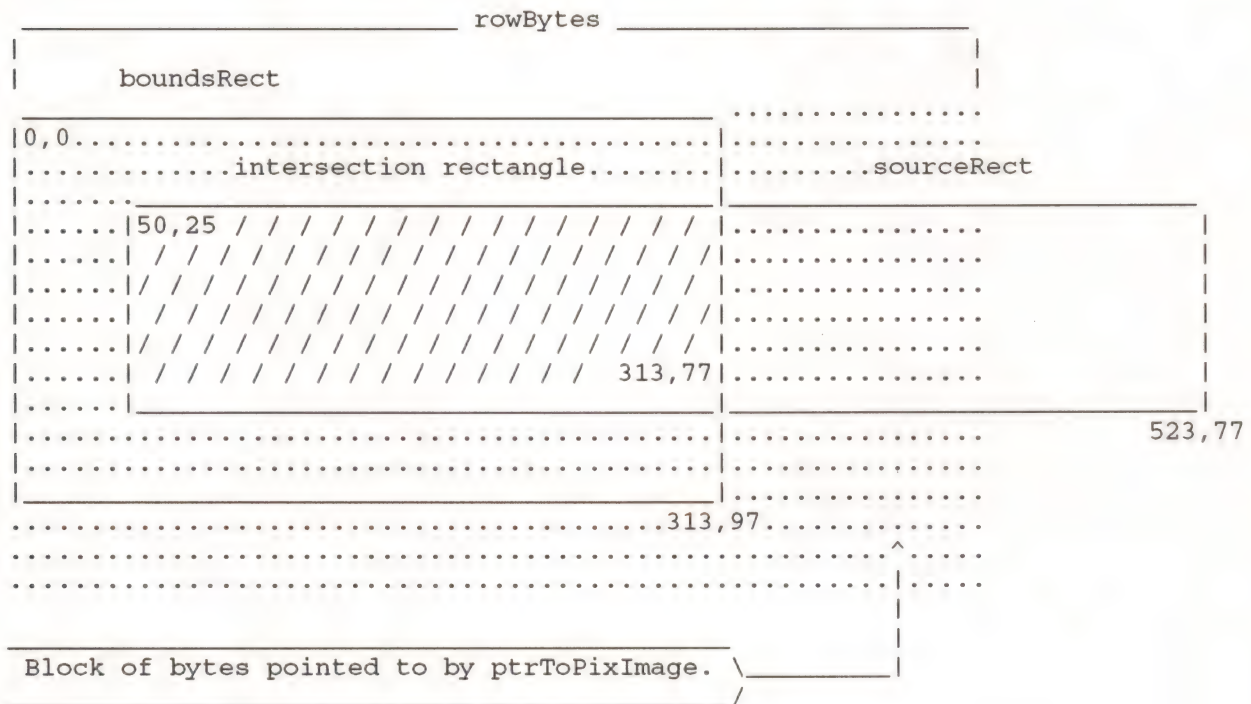


Figure 2-Sample LocInfo Structure With sourceRect

Figure 2 contains a sourceRect which is not completely contained by the boundsRect; the sourceRect is so wide that it even goes beyond the edge of the

pixImage. If the entire contents of this rectangle were drawn, the result would be quite a mess, since it extends beyond the boundary of the pixelmap. However, PPToPort first intersects the sourceRect and the boundsRect, and then uses the resulting intersection rectangle (illustrated with a thicker border in the figure). PPToPort uses only the contents of the intersection rectangle.

Up until now, the boundsRect upper-left corner has always been 0,0. This is an easy way to think of it, but it is not necessary. The important thing to remember about these rectangles is their relation to one another. If one were to offset both the boundsRect and sourceRect in this example, the values for the corners of the rectangles would change, but the relationship between the two rectangles would stay the same. Figure 3 illustrates the same example if one were to offset both rectangles by -60,-45.




Figure 3-Sample LocInfo Structure Offset by -60,-45

Notice that the same area of the pixImage is involved, even though the boundsRect and sourceRect are offset. When one offsets both the boundsRect and sourceRect by the same amount, the referenced part of the pixImage does not change--this is an important concept.

Time to ask a question that is answered shortly: "Why isn't the upper-left corner of the boundsRect always 0,0?" Because the LocInfo record isn't always a source LocInfo record. It can also be a destination LocInfo record, and the most common pixelmap to which a destination LocInfo record refers is the screen.

If you had not noticed, the discussion changes gears here--to discuss LocInfo records that indicate a destination pixelmap. Basically, everything is the same as has been described with two exceptions. First, destination pixelmaps do not have a sourceRect. Instead there is a rectangle that describes some portion of the destination pixelmap, and this rectangle is called the portRect. Second, the LocInfo record is part of a grafPort, and each grafPort has a LocInfo record as part of the grafPort data structure.

It is important to remember that a LocInfo record can be used as either a

source or destination LocInfo. All a LocInfo record does is define some bytes in memory as a pixImage. Even the screen, which is usually used as a destination pixelmap, can be used as a source pixelmap. There could be situations where one might want to take part of the screen and copy it into some off-screen pixelmap, and in this case, the screen would be a source of pixel data, not a destination.

In the case of the screen pixelmap, there are no wasted bytes in the pixImage, as all of the screen bytes are enclosed by the boundsRect. The screen width of 160 is evenly divisible by eight, so there is no slop at the right edge, and there are no extra rows hanging off the bottom of the boundsRect.

Figure 4 shows a sample LocInfo and portRect (every grafPort has a LocInfo and a portRect).




Figure 4-Sample LocInfo and portRect

Following are two important points to remember:

1. Every grafPort works in local (not global) coordinates (local coordinates are defined soon).
2. The origin of the grafPort is the upper-left corner of the portRect. There is no GetOrigin call; there is a SetOrigin call, but no GetOrigin. To get the origin of a grafPort, one needs to do a GetPortRect call, and then look at the upper-left corner to determine the current origin of the grafPort. This is the way to get the origin.

In the case of Figure 4, local and global coordinate systems are the same, as is always the case when the boundsRect has an upper-left corner of 0,0 (which it seldom does). So, for this exceptional case, one doesn't need a definition of local coordinates. In the global coordinate system, the upper-left corner of the screen is 0,0. In local coordinates, the upper-left corner of the screen is whatever the boundsRect says it is. So when the upper-left corner of the boundsRect is 0,0, the global and local coordinate systems are the same.

In Figure 4, if one tried to draw something to point 0,0, it would not draw--it

excluded part is clipped to protect the window above from being drawn upon. This is how window clipping works. This is all there is to it.

Figure 6 enhances Figure 5 by adding an overlapping window to demonstrate the visRgn.

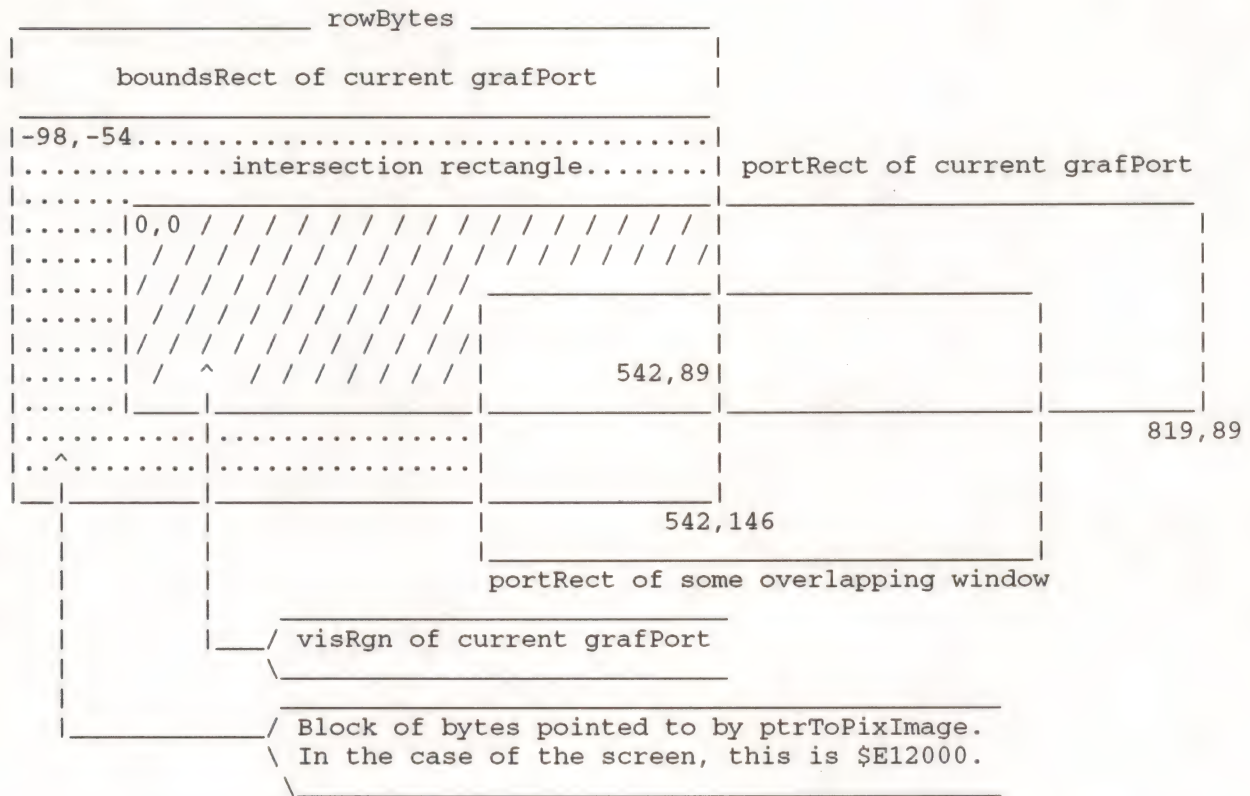


Figure 6-Sample LocInfo and portRect With Overlapping Window

What happens to the visRgn during a SetOrigin? Remember that the boundsRect and portRect get offset. The visRgn does too. Again, if all of these elements are offset together, then the relationship between them remains the same; they stay the same, relative to one another. (For more information, see Einstein's theory of general relativity.)

The final component for clipping is the clipRgn, which is the application's property and, therefore, the application's responsibility. The system sets the clipRgn about as big as it can get to start (much bigger than the portRect); this is often referred to as arbitrarily large, even though it isn't so arbitrary. The system creates all grafPort structures with a large clipRgn, and this can be a problem for certain types of QuickDraw II operations. Since the clipRgn already reaches to the borders of the conceptual drawing space, it cannot be offset; it is effectively stuck, due to its size. It is a good practice to make the clipRgn smaller than the system default.

SetOrigin does not offset the clipRgn. (This is why the size problem with a big clipRgn is not so apparent.) The clipRgn is the only clipping component that is not offset by SetOrigin, and one should consider this when using clipRgn for clipping effects, since an application must remember to offset it if it needs to be offset.

Now with all of the fundamentals out of the way, it is time to play some grafPort clipping games. As a refresher, there are four clipping components in a grafPort: the boundsRect, the portRect, the visRgn, and the clipRgn.

If an application creates its own off-screen grafPort structures, then it can do as it wishes with all four clipping components. After all, if it has the responsibility to set them up in the first place, it should have the right to change them. If, however, the Window Manager creates the grafPort structures, then an application should keep its figurative hands off certain clipping components, namely the boundsRect and the visRgn. The clipRgn, by definition, is the application's to do with as it sees fit, and if careful, an application can also change the portRect. Changing the portRect can be very useful, but one needs to be careful and fully understand all of the ramifications.

So, why would one change the portRect, and how would one do it?

Another figure is in order.




Figure 7-Sample LocInfo and Modified portRect

One can use the GetPortRect call to get the portRect for the current grafPort. One can then modify it, and then use the SetPortRect call to inform the grafPort about the change. Why do this? In Figure 7, the dotted line represents the new left edge of the portRect after the modification (a simple modification of adding 100 to the old value of zero).

Note that changing the portRect in this way changes the relationship between the portRect and the boundsRect. Anything drawn from 0 to 99 (x coordinate) is clipped, since it is outside the new (modified) portRect. Before the modification, anything drawn from 0 to 99 would have affected the screen.

This modification may cause the portRect to be smaller than the visRgn. This is okay, since the visRgn can only cause more clipping, not less. So, all of this works just fine. Note that the origin changed when the left edge of the portRect changed. The upper-left corner of the portRect is always the origin, and an application changed it. The origin changed without a SetOrigin call. (Scary, huh?)

One could have done exactly the same thing by making a clipRgn to exclude the x coordinates from 0 to 99. However, here is something cool. After the modification, do a SetOrigin(0,0), which sets the upper-left corner of the

shrunk portRect to 0,0. One cannot accomplish this sort of thing as simply by making a clipRgn. One can effectively move where an origin of 0,0 is the screen, and just building a clipRgn to exclude some part of the screen does not accomplish this.

Why would one want to change where 0,0 is on the screen? This sort of trick is very useful for adding rulers to a document window, for example. One of the problems with rulers is that they should not scroll with the rest of a document. Unfortunately, TaskMaster, if allowed to handle scrolling, doesn't know about a ruler at the top of a window and scrolls it with the rest of the window's content area. By changing the portRect so that the ruler is not inside of it, one can keep TaskMaster from scrolling it. In a draw procedure, when it is necessary to draw the ruler, grow the portRect, set the origin to 0,0, and then draw the ruler. Once it is drawn, set the portRect back to the smaller size to protect the ruler again.

Another reason one might want to do this is if an application uses a split window (where the top of the window may show a different part of the document than the bottom). Changing the portRect has the advantage that the upper-left corner of the portRect is always the origin, so it makes mapping document coordinates easier.

Another advantage to using the portRect in this way is that it keeps the clipRgn free for other purposes. Being able to separate types of clipping to either the portRect or the clipRgn keeps the clipRgn from being overused.

As a final note, it should be observed that the only clipping that is done is on a destination pixelmap. There is no clipping on a source pixelmap. There is no need. All the clipping needed is done at the destination end, so it would be wasteful to clip twice.

This finishes the discussion about QuickDraw II and how the boundsRect, portRect, visRgn, and clipRgn work together to accomplish clipping. Hopefully this Note answers more questions than it creates.

Further Reference

-
- o Apple IIGS Toolbox Reference, Volume 2
 - o Relativity the Special and General Theory (1920)

Apple IIgs
#81: Extended Control Ecstasy

Revised by: Dave Lyons
Written by: C.K. Haun

July 1991
May 1990

This Technical Note discusses special features and concerns that should be considered when using the extended controls introduced in System Software 5.0.

Changes since November 1990: Added information on which fields the Control Manager automatically copies from a custom control's template to its control record. Corrected NewControl2 parameter order. Added a note about SetCtlTitle. Changed some "pea 0" instructions into "pha" when pushing space for results.

Introduction

The extended controls introduced in System Software 5.0 allow the application programmer a great deal more freedom in designing and controlling applications. The new features enhance the functionality of the controls and TaskMaster, but can cause confusion and consternation if you are careless with the new parameter block. This Note also includes a discussion of the multipart nature of many of the extended controls and some pointers for writing extended custom controls.

Counting The Costs

One of the major stumbling blocks seen when programming the new extended controls is bad parameter counts. Extended controls introduce parameter blocks and parameter counts to the Control Manager. You need to fully understand the parameters required and the resulting parameter count for each control you create, or you can experience program problems that may be very confusing and difficult to track down.

Remember also that the Control Manager does not understand "skipping parameters." If you are creating an extended radio button and you want key equivalents, but not a color table, you cannot ignore the color table parameter field. There is no way to tell the Control Manager to "skip" a field during control creation; make sure you initialize all the fields to values that are meaningful--either a real pointer, handle, resource ID, or zeroes. In this case, if you try to "skip" the color table and do not zero out the color table parameter field, the resulting radio button wears ugly colors you do not expect.

As you can imagine, miscounting parameters in other extended controls can produce confusing results, so the parameter count is the first place you should check when you're having difficulties creating extended controls.

For Rez users, the Types.Rez file contains extended control templates that automatically generate the correct count value, so programmers creating all their controls with the Resource Compiler should not have these problems.

Silly Little Bits

The other area of the new parameter block model that is giving folks trouble is the moreflags field. This field hones the definition given by the reference fields and can cause you much grief if misused. Make sure to set the reference bits to the values you require. The bit settings have been standardized across all the extended controls, with %00 indicating a pointer, %01 indicating a handle, and %10 indicating a resource. Remember also to set the bits for all the references you have--strings, color tables, or whatever else may be ambiguously referenced in the control.

If you accidentally use the wrong bit pattern, you can experience strange bugs ranging from garbled text to SysFailMgr caused by a nonexistent resource being referenced. Again, Rez users can use the equates for all the reference specifiers in the Types.Rez file to avoid confusion and evil bugs.

The Parts are Greater than the Whole

To create some new extended controls, like pop-up menu controls and LineEdit controls, functions of different tool sets were combined. Pop-up menu controls are a combination of the Control Manager and the Menu Manager, LineEdit controls are a blending of the Control Manager and the LineEdit tool set, and other new control types follow the same pattern.

This means that, at times, you have to go further into the documentation to find information. Getting the text out of an LineEdit extended control is a multistep process that is a good example of this type of problem.

```

MyLineEdit dc    i2'8'          ; parameter count
            dc    i4'1'          ; id number 1
            dc    i2'10,10,23,90' ; control rectangle
            dc    i4'editLineControl' ; process reference
            dc    i2'0'          ; flags
            dc    i2'fCtlCanBeTarget+fCtlWantEvents+fCtlProcRefNotPtr'
                                ; moreflags
            dc    i2'0'          ; refcon
            dc    i2'15'         ; maximum characters allowed
            dc    i4'0'          ; no default text
MyLineEditHandle ds 4          ; handle for the created control

        pha
        pha
        pushlong mywindowgrafport ;window that control will residein
        pea    0                  ;verb, single Extended control
        pushlong #MyLineEdit
        _NewControl2
        pla
        sta    MyLineEditHandle   ; save the control handle
        pla
        sta    MyLineEditHandle+2

```

When you want to get the text back out of that control later, you begin to experience what it means to have a control that is an amalgam of various tools. You would start by using the control handle returned by NewControl2:

```

Scratch    equ    $0              ; some scratch space
Scratch2   equ    $4

        lda    MyLineEditHandle   ; move the ControlHandle to
        sta    Scratch            ; some direct page space
        lda    MyLineEditHandle+2
        sta    Scratch+2

```

```

lda    [Scratch]
tax
ldy    #2
lda    [Scratch],y          ; and dereference it, putting it
                               ; back in some dpage
sta    Scratch+2          ; space to use it.
stx    Scratch

```

That gives you the pointer to the control record. Stored in the control record is the handle of the LineEdit item that is actually controlling the text processing:

```

pha                    ; make space for the text handle to
                       ; be returned
pha
ldy    #oct1Data       ; offset to the ctlData section
                               ; of the ControlRecord
lda    [Scratch],y    ; where the handle for the actual
tax
iny
iny
lda    [Scratch],y
pha
phx
_LEGetTextHand        ; ask for the handle for the text
pla                    ; in this LineEdit control
sta    Scratch2        ; and now you have the handle to
                       ; the text you want.

pla
sta    Scratch2+2

```

The main point is that when you are using extended controls, you often cannot use the Control Manager to do everything that needs to be done. You also need to understand and use the supplementary or "hidden" tool sets.

Here's another example, using a pop-up menu extended control, and in this case we define a font pop-up that contains all the font names currently available.

```

MyPopUpControl dc i2'9'          ; parameter count of 9
                dc i4'1'          ; control ID of 1
                dc i2'2,2,0,0'    ; Position, upper left corner of
                               ; the window, let
                               ; Control Manager calculate full
                               ; size
                dc i4'popUpControl' ; def proc for PopUp
                dc i2'0'          ; flags
                dc i2'fCtlWantEvents+fCtlProcRefNotPtr'
                               ; more flags
                dc i4'0'          ; ref con
                dc i2'0'          ; title width, will be calculated
                dc i4'mymenu'     ; pointer to actual menu structure
                dc i2'500'        ; initial value, item number of
                               ; item to be displayed in popup at
                               ; creation

mymenu          dc i2'0'          ; version number, should be 0
                dc i2'200'        ; menu ID number
                dc i2'0'          ; menu flags
                dc i4'mymenutitle' ; pointer to menu title
                dc i4'mymenuitem1' ; first menu item

```

```

        dc    i4'mymenuitem2'          ; second menu item
        dc    i4'0'                    ; null terminator, end of menu
mymenutitle str 'Font'

mymenuitem1 dc i2'0'                    ; version number
            dc i2'500'                  ; item number
            dc i2'0'                    ; no hot keys
            dc i2'0'                    ; not checked
            dc i2'0'                    ; item flags, no special drawing
            dc i4'mymenuitem1title'
mymenuitem1title str 'Plain'

mymenuitem2 dc i2'0'                    ; version number
            dc i2'501'                  ; item number
            dc i2'0'                    ; no hot keys
            dc i2'0'                    ; not checked
            dc i2'1'                    ; item flags, bold face this one
            dc i4'mymenuitem2title'
mymenuitem2title str 'Bold'

```

Now create this control:

```

pha
pha
pushlong mywindow                      ; target window grafptr
pea    0                                ; verb, single control pointer
pushlong #MyPopUpMenu
_NewControl2
pullong mypopuphandle                  ; save the handle

```

This pop-up menu control created is not associated with the menu bar across the top of the desktop. You can consider each of your pop-up menu controls as separate menu bars, so if you want to perform Menu Manager calls on a pop-up menu control, you need to set the menu to point at your pop-up menu control. In this example, to add all the fonts available to the pop-up menu you would:

```

pha
pha                                     ; space to hold current bar
_GetMenuBar                             ; get the handle to the current
                                         ; menu bar

pushlong mypopuphandle
_SetMenuBar
pea    200                               ; id number of this menu
pea    502                               ; first font family ID number to
                                         ; use
pea    0                                 ; fontspecbits
_FixFontMenu
pea    0
pea    0
pea    200
_CalcMenuSize                           ; re-size the popup menu
_SetMenuBar                             ; restore the previous menu as the
                                         ; current menu

```

Controls That Are Not Controls

The new picture extended control is not a "full-fledged" control; it has been provided to simplify your programming tasks. The picture control does not support normal mouse hit testing and highlighting. Think of it as a built-in extension to your content drawing routine, and not as a control. It is

provided to allow you to refresh your whole window with a single DrawControls call, instead of drawing the controls and then drawing pictures. The icon button extended control has been provided as the graphic full-function control. If you need or want a fully functional control that uses a picture, you should consider writing your own custom control procedure.

Custom Extended Controls

Custom controls can also benefit from all the advantages of extended controls. You can create a custom control that uses a template, can be a resource, has a definition procedure that is a resource, and responds to all the new control calls. If you write an extended custom control or upgrade a previously-written custom control, there are new messages and changes to existing messages of which you need to be aware. These changes are documented in volume 3 of the Apple IIgs Toolbox Reference.

The Control Manager copies the following fields from the control template to the control record before it sends your control the init message: `ctlOwner`, `ctlID`, `ctlRect`, `ctlFlag`, `ctlHilite`, `ctlMoreFlags`, `ctlVersion`, `ctlRefCon`, and `ctlProc`. The `ctlNext` field is owned by the Control Manager. If any additional fields need to be set up based on the control template (such as `ctlValue`, `ctlData`, `ctlColor`, and any custom fields), your init routine needs to take care of it.

Putting your custom control definition procedure in a resource can significantly enhance the functionality of the custom control. You may find it easier to add to all of your programs and you do not have to manage the code space required. If you do write a custom control definition procedure and want to store it as a resource, here are some hints for success.

First, the code you store in your resource fork must be fully compiled and linked code. The code resource converter uses the System Loader to load the code, so the code must be executable code, not object code.

Second, set the convert and locked bits of the resource attributes for your code resource. The convert bit must be set to tell the Resource Manager to call the code resource converter when it loads this resource. The resource type for control definition procedures is `rCtlDefProc`, `$800C`.

By setting locked but not fixed, memory fragmentation is reduced (because of how the code resource converter and Memory Manager work). Setting the locked attribute is also recommended for compatibility with future system software.

Third, keep in mind that this definition procedure may be purged and reloaded whenever the Memory Manager needs the space. This means that you cannot store any information in your definition procedure if you want to keep track of it between calls to the definition procedure. If you do, and your definition procedure gets purged and reloaded, you lose that data.

If you need data space for your custom control, use the control record as your stash. You can easily either use the fields already provided in the control record, or you can expand the control record to as much space as you need (within sensible limits) and store your data there.

Warning: Control definition procedures are initially loaded with purge level zero. When they are released, they are given purge level three. If they are then reloaded, the Resource Manager does not change the purge level back to zero--your definition procedure may then be purged (even while executing) unless its handle is locked. The solution is to lock your definition procedure handle within the procedure:

```
myPosition    pea    0                ; space for result
```

```
pea    0
pushLong #myPosition
_FindHandle
_HLock
```

and unlock your handle with HUnlock on exit. This keeps your procedure safe, while not creating "code islands," which clog up memory.

Changing a Control's Title

If you call SetCtlTitle to give a control a new title, everything is great if the new title is referenced the same way as the current title (by pointer, by handle, or by resource ID). If the new title is referenced differently, you must first call SetCtlMoreFlags on your control so that the SetCtlTitle value can be interpreted correctly.

Conclusion

The extended controls provided in System Software 5.0 and later are a great leap forward for programmers. They relieve the application of much of the tedious detail code that relates to housekeeping, not the guts of application programming. Used in combination with the enhanced TaskMaster, you can have an application's visual interface up and running a lot faster, leaving you more time to work on the heart of your application.

Further Reference

-
- o Apple IIgs Toolbox Reference, Volumes 1-3.

Apple IIgs

#82: Controlling the Control Manager

Revised by: Matt Deatherage

November 1990

Written by: Dave Lyons

May 1990

This Technical Note describes an anomaly in the NewControl2 call in System Software 5.0.2 and provides a solution.

Changes since May 1990: Noted that System Software 5.0.3 fixes this anomaly.

This Note formerly advised of a problem with NewControl2-the current port was not set before adding the controls, which gave unpredictable results. System Software 5.0.3 and later fix this problem.

Apple IIgs
#83: Resource Manager Stuff

Revised by: Matt "Even less of a middle name" Deatherage
Written by: Dave Lyons

May 1992
May 1990

This Technical Note answers your miscellaneous Resource Manager questions.

CHANGES SINCE DECEMBER 1991: Added several notes pertaining to System Software 6.0 and a note about making Resource Manager calls from a resource converter. Added new discussion about how "changed" is really a resource attribute.

UNIQUERESOURCEID

In System Software 5.0.4 and earlier, calling UniqueResourceID with an IDRange value of \$FFFF does not work reliably. It sometimes returns a system-range ID (\$07FFxxxx) if there are already system-range resources of the specified type present in the current search path.

If you are using a development utility that generates resource IDs using UniqueResourceID, check the results to make sure no system-range resource IDs are being used by accident. This problem is fixed in System Software 6.0.

WHAT SETCURRESOURCEFILE DOES

SetCurResourceFile is documented in Chapter 45 of the Apple IIgs Toolbox Reference, Volume 3 (see especially "Resource File Search Sequence" near the beginning of the chapter).

This explanation might make you think SetCurResourceFile rearranges the search path, but it does not; instead, it just makes searches start at a different place in the path. SetCurResourceFile is useful for controlling what resource files are searched, not for changing the search order.

HOW THE TOOLBOX USES RESOURCES AS TEMPLATES

The toolbox uses several types of resources as templates for creating other objects. Examples include rControlList, rControlTemplate, and rWindParam1. The toolbox automatically releases these resources from memory as soon as it is through with them, so there is no need to create your template resources with special purge levels in an effort to free more memory. It is not a problem.

USING RESOURCES FROM WINDOW UPDATE ROUTINES

In System Software 6.0 and earlier there is no special code to set the current resource application when the system calls an application window update routine (See Apple IIgs Technical Note #71 for notes on NDAs and the current resource application).

To avoid a situation where a window update routine cannot get needed resources, obey the following rules:

1. Application window update routines must either (a) assume that the resource application has the same value it had when the window was created, or (b) save, set, and restore the current resource application, using `GetCurResourceApp` and `SetCurResourceApp`.
2. NDAs that start the Resource Manager must not call application window update routines, and they must not cause application window update routines to be called (for example, if an NDA calls `TaskMaster` to handle a modal dialog or movable modal dialog, the `tmUpdate` bit in `wmTaskMask` must be off).

CURRESOURCEAPP IN INFODEFPROCS AND CUSTOM WINDOWS

The current resource application has no guaranteed value when an information bar definition procedure or custom window definition procedure gets control. These must always save, set, and restore the current resource application using `GetCurResourceApp` and `SetCurResourceApp`.

STARTUPTOOLS OPENS RESOURCE FORKS READ-ONLY

When `StartUpTools` opens your application's resource fork, by default it opens it with read-only access. If your application needs to make changes to the resources on disk in System Software 5.0.4 and earlier, you need to close the fork and reopen it with read and write access. To close it, use `GetCurResourceFile` and `CloseResourceFile`; to reopen it, use `LGetPathname2` and `OpenResourceFile`.

Note : You must update the `resFileID` field in the `StartStop` record if you close and reopen your resource fork. `CloseResourceFile` disposes the handles of any resources in memory from the file you're closing, so you must call `DetachResource` on any resources you need to keep. (If you pass an `rToolStartup` resource to `StartUpTools`, the system detaches it for you automatically.)

In System Software 6.0 and later, setting bit 3 (`$0008`) of the `startStopRefDesc` tells the Tool Locator to open your resource fork with all allowed permissions instead of with just read permission.

CALLING STARTUPTOOLS FROM A SHELL APPLICATION (FILE TYPE \$B5, EXE)

In System Software 5.0.4 and earlier, `StartUpTools` tries to open the current application's resource fork. It determines the pathname of the "current application" by examining prefix 9: and making a `GET_NAME` GS/OS call, but do not assume it will always construct the pathname this way. If you call `StartUpTools` from a shell application and expect it to open your EXE file's resource fork, you will be disappointed.

If GS/OS has launched your application, life is good--usually, though, a shell has loaded your shell application directly, so `GET_NAME` returns the name of the shell instead of the name of your application file.

To open your shell file's resource fork, call `ResourceStartUp`, get the pathname by calling `LGetPathname2` on your user ID, and pass the pathname to `OpenResourceFile`. `StartUpTools` uses this strategy all the time in System Software 6.0 and later, meaning you don't have to.

WHAT'S NIL IN A RESOURCE MAP?

The resource maps for open resource files are kept in memory, and the structure is defined in chapter 45 of Apple IIgs Toolbox Reference, Volume 3.

The `resHandle` field of a resource reference record (`ResRefRec`) is defined as "Handle of resource in memory. A NIL value indicates that the resource has not been loaded into memory." In this case, NIL means that the middle two bytes of the four-byte field are zero. In other words, a NIL entry in the resource map may have a non-zero value in the low-order byte.

LOADRESOURCE AND SETRESLOAD(FALSE)

When you call `LoadResource` on a locked or fixed resource and `SetResLoad` is set to `FALSE`, you may get Memory Manager error \$0204 (`lockErr`), because the Resource Manager tries to allocate a locked or fixed zero-length handle, which the Memory Manager does not permit.

ADJUSTING THE SEARCH DEPTH

If you wish to add some resource files to the beginning of a resource search path and adjust the depth so that the end point of the search is unchanged, it's tempting to use `SetResourceFileDepth(0)` to get the current depth, add one, and set this new depth with `SetResourceFileDepth`.

The problem is that the search depth is often `-1 ($FFFF)`, meaning "search until the end of the chain." If you add your adjustment to `-1`, you do not usually get the intended effect. A solution is just to check for `$FFFF` and not adjust the depth in that case.

CURRESOURCEAPP AFTER RESOURCESHUTDOWN

After a `ResourceShutDown` call, the current resource application is always `$401E`. (The Resource Manager starts itself up at boot time with its own memory ID, `$401E`. Do not ever call `ResourceShutDown` while the current resource application is `$401E`.)

RESTORING THE CURRESOURCEAPP

If you need to start up and shut down the Resource Manager without disturbing the current resource application, call `GetCurResourceApp` before `ResourceStartUp`, and call `SetCurResourceApp` to restore the old value after `ResourceShutDown`.

It does not help to call `GetCurResourceApp` after `ResourceStartUp`, since the application just started up is always the current resource application.

Shell programs which start the Resource Manager need to call `SetCurResourceApp` after regaining control from a subprogram (for example, an EXE file) which may have started and shut down the Resource Manager, leaving the current resource application set to `$401E` instead of the shell's ID.

Shell programs that do not start the Resource Manager have nothing to worry about. In this case the current resource application is normally `$401E`, so when a subprogram calls `ResourceShutDown` life is still wonderful.

WHAT INFORMATION IS KEPT FOR EACH RESOURCE APPLICATION?

When you switch resource applications with `SetCurResourceApp`, that takes care

of all the application-specific information the Resource Manager has.

There is no need to separately preserve the current resource file, the search depth, the SetResourceLoad setting, or any application resource converters that are logged in. All of this information is already recorded separately for each resource application.

"CHANGED" IS A RESOURCE ATTRIBUTE

This seems obvious when first reading the documentation, but it has a consequence that isn't so obvious.

If you mark a resource as changed with MarkResourceChanged and later use SetResourceAttr to change that resource's attributes, you must include resChanged in the attributes you specify or the Resource Manager does not still know the resource has changed.

This means you can undo a MarkResourceChanged call, but it also means you need to preserve the resChanged bit across SetResourceAttr calls if you don't want to accidentally achieve the same effect.

The Resource Manager clears the resChanged attribute when a resource is written to disk; the attribute indicates the data in memory is more recent than what's on disk. Normally, adding a resource with AddResource sets this bit because the resource isn't actually written to disk until the resource file is updated.

However, if AddResource has to make the file longer (by extending the EOF), it writes the resource to disk immediately. This means that in some cases, a resource added with AddResource will be properly added but the resChanged attribute will not be set. Don't be confused if this happens to you.

MAKING RESOURCE MANAGER CALLS FROM RESOURCE CONVERTERS

Don't. This would be a first-class example of reentrancy, and the Resource Manager is not reentrant in any class.

WHO OWNS HANDLES PASSED TO ADDRESOURCE?

When you pass a handle to AddResource, the Resource Manager is responsible for the handle unless AddResource returns an error. Once you call AddResource, the handle belongs to the Resource Manager and you must treat it like you would the handle to any other resource.

NAMED RESOURCE BUGS IN SYSTEM SOFTWARE 6.0

The new-for-6.0 Resource Manager function RMFindNamedResource compares the resource name you requested to named resources incorrectly. The comparison algorithm doesn't compare the lengths of the strings before starting to compare the characters. This means, for example, that if you request a resource named "Raymond" and the Resource Manager encounters a named resource named "Raymond" first, it will return the resource named "Raymond" instead. This anomaly also affects the HyperCard IIgs named-resource XCMD callback functions, even though they don't use the Resource Manager's named-resource calls.

This anomaly also affects RMLoadNamedResource, which calls RMFindNamedResource.

DEBUGGING INFORMATION

The following information is provided for your convenience during program development. It allows you to check exactly what user IDs are using the Resource Manager, what files are in their search paths, and what resource converters are logged in.

DO NOT depend on this information in your program; it is subject to change in future versions of the Resource Manager.

All the Resource Manager's data structures are rooted in the Resource Manager tool set's Work Area Pointer (WAP). To get the Resource Manager's WAP, call GetWAP (in the Tool Locator) with userOrSystem = \$0000 and tsNum = \$001E.

The WAP value is a handle to the Resource Manager's block of global data. Several interesting areas in this block are listed below.

+ \$0A2	curApp	Word	Offset into the globals block of the current resource application's Application Record.
+ \$2B0	sysFile	Long	Handle of system file map, or NIL if none.
+ \$2B4	sysConvertList	Long	Handle of system converter list, or NIL if none.
+ \$2B8	appList	20*n bytes	List of Application Records (20 bytes each).

Each Application Record has this format:

+000	appFlag	Word	0=entry available, 1=entry used, \$FFFF = end of array.
+002	appID	Word	User ID of application.
+004	appFiles	Long	Handle of application's first resource map, NIL=none.
+008	appCur	Long	Handle of application's current resource map, NIL=none.
+012	appConverters	Long	Handle of application's converter list, NIL=none.
+016	appReadFlag	Word	1=read resources, 0=don't read (SetResourceLoad).
+018	appFileDepth	Word	Number of files to search in this path.

Converter lists have this format:

+000	n	Word	Number of entries in the table (entries can be unused).
+002	theConverters	6*n bytes	List of converter entries (6 bytes each).

Each Converter entry has this format:

+000	resType	Word	Resource type for this converter (\$0000 for unused entry).
+002	convAddress	Long	Address of resource converter.

The format for a resource map is described starting on page 45-17 of Apple IIgs Toolbox Reference, Volume 3.

Remember, don't depend on this information in your application; use it during debugging, and use it to write debugging utilities.

Further Reference

- o Apple IIgs Toolbox Reference, Volume 3
- o Apple IIgs Technical Note #71, DA Tips and Techniques

Apple IIgs
#84: TaskMaster Madness

Written by: C.K. Haun <TR>

July 1990

This Technical Note discusses the enhancements made to TaskMaster in System Software 5.0.

TaskMaster has been expanded to handle extended control actions and give you more information about events in System Software 5.0. This Note discusses some features of the expanded TaskMaster and TaskMasterDA, and how you can best exploit the new features in your applications.

Stop Making It So Difficult

Developers just want to work too hard. You get a neat new thing like the expanded TaskMaster, and you still want to do all the work yourself. The new TaskMaster does nearly everything for you, as long as you treat it correctly.

What this means is you do not have to call FindControl, TrackControl, TEIdle, LEKey, handle keystrokes for controls, keep track of click counts, or any of the other mundane event management tasks unless you specifically want to perform actions that TaskMaster does not perform. For the standard controls and situations this means that you do not have to do anything.

The magic keys to this life of freedom and ease are the five newly defined taskMask flag bits, labeled in the interfaces as tmContentControls, tmControlKey, tmControlMenu, tmMultiClick and tmIdleEvents. This Note looks at what the new bits do for you, but first a word of warning.

Warning: If you set any of these new bits, TaskMaster assumes you are using the new extended task record. This means that you cannot just go into an older program and set these bits and expect your program to work successfully. You also must allocate the additional space for the extended portion of the task record. If you do not, TaskMaster puts task data in areas that you do not expect, and Bad Things happen.

Bits 'o This, Bits 'o That

Click Bits

tmMultiClick tells TaskMaster to keep the new "click information" fields in the extended task record updated. This allows you to have TaskMaster keep track of multiclick events; the wmClickCount field is one, two or three depending on whether the last action was a single, double, or triple click. In fact, if you can click your mouse button fast enough, you can time quadruple clicks, sextuple clicks, or as high as you want, although anything over triple-clicking is nearly impossible for users to consistently manage. wmClickCount just gets incremented by one when the click falls within the double time interval. wmLastClickTick is updated with the system tick value at last click. wmLastClickPt contains the location of the last mouse click. TaskMaster calls GetDbtTime internally to determine the correct time intervals for these values.

Idle Bits

`tmIdleEvents` tells `TaskMaster` to call the idle routines for controls that need idle events, like `TextEdit` controls and `LineEdit` controls. This also means that only the active control is blinking a cursor, since `TaskMaster` is working with the target bits of the extended control records to keep track of which `TextEdit` or `LineEdit` control is active and switching the target control in response to mouse clicks and Tab keypresses. This is also the area where you tell `TaskMaster` how to highlight your window controls. Using the `Control Manager` calls `MakeNextCtlTarget` and `MakeThisCtlTarget` allows you to specify which `LineEdit` or `TextEdit` control is active. You can use these calls to highlight input errors the user has made. For example, if someone has entered text in a `LineEdit` control that requires a number, you can alert the user if he enters non-numeric characters with an `Alert` or `AlertWindow` call. You can then direct the user to the `LineEdit` control that contains the bad entry by calling `MakeThisCtlTarget` with the handle of that `LineEdit` control. This deactivates any other target control and moves the insertion point to the `LineEdit` control that needs the correction.

Contentious Bits

`tmContentControls`, `tmControlMenu` and `tmControlKey` bits are the real workhorses of the expanded `TaskMaster`.

When the `tmContentControls` and `tmControlMenu` bits are set, `TaskMaster` handles the mouse activity side of events--tracking, highlighting or popping-up the selected control. If the control is a radio button, check box, pop-up menu or list control, `TaskMaster` also performs the correct action for the click, either setting the control value, scrolling the list, setting the pop-up menu to the selected item, and so on. `TaskMaster` then returns a `taskCode` of `wInControl` (\$21). The control handle is stored in `wmTaskData2`, the part code of the part selected in `wmTaskData3` and the control ID is in `wmTaskData4`. For many of the controls in your windows your application needs to take no further actions, `TaskMaster` has set the control values. When the user closes the window or clicks on a button that causes an action, you can then read the values of all the controls you care about at that point and do what you need to do, instead of keeping track as the user manipulates controls.

The last new bit, `tmControlKey`, works with the `tmControlMenu` bit to handle key events for your extended controls.

When a key event occurs, `TaskMaster` sends the event to the internal routine `TaskMasterKey`. `TaskMasterKey` first looks at the `tmMenuKey` bit (which has been in `TaskMaster` since the `Window Manager` was implemented). If it is set, then `TaskMaster` tries to handle the event as a menu event, calling `MenuKey` for the current menu bar.

Note: This also means that any key equivalents in your main menu bar (across the top of the desktop) take precedence over key equivalents in your window controls.

If this fails (or that bit is not set) and `tmControlKey` is set, then `TaskMasterKey` polls the controls in the currently open window for any controls that would like this keystroke, either for controls with a `keyEquivalent` field or a pop-up menu control with key equivalents for menu items. If it finds a control that wants the key event, it is handled very much like a mouse event. The action for the control is performed (checking a check box, for example) and the `wmTaskData` fields are filled as they would be for a mouse click, and an event code of `wInControl` (\$21) is returned. If a key event did occur, you can differentiate it from a mouse event by looking at the `wmWhat` field of the `taskRecord`. Even though a `wInControl` event code was passed back by

TaskMaster, the `wmWhat` field is either `$0001` or `$0003`, the former for a mouse down event and the latter if a keystroke stimulated the `wInControl` event.

Even More Bits

All these new features rely very heavily on the changes made to the Control Manager in System Software 5.0. Many of the TaskMaster features, keystrokes, target controls, and so on only work if you have the `moreFlags` bits set correctly in your control definitions. If you are having difficulty with new TaskMaster features, check your control definitions against the information in the Control Manager chapter of Volume 3 of the Apple IIgs Toolbox Reference and Apple IIgs Technical Note #81, Extended Control Ecstasy.

Don't Get Goofy

There are some dangers in these new features, of course. By allowing built-in key equivalencies for almost all the controls that can exist in a window, it may be tempting to define key equivalents for everything, and create weird and unusual key combinations for your controls. Please remember the Human Interface Guidelines (specifically Human Interface Note #8, Keyboard Equivalents) and keep your use of keystroke equivalents to a minimum. Multimodifier keystrokes (Command-Option-Shift, for example) do not enhance the user's experience and can be very confusing.

NDA's Can Have Fun Too

TaskMasterDA has also been added to the Window Manager, providing your new desk accessories (NDAs) with the same kind of TaskMaster support your applications have. This lets you easily use extended controls inside NDAs, following the same basic rules as in an application. There are only a few things to worry about.

What Does That Stack Picture Really Mean?

The input to TaskMasterDA, as shown in Volume 3 of the Apple IIgs Toolbox Reference, is as follows:

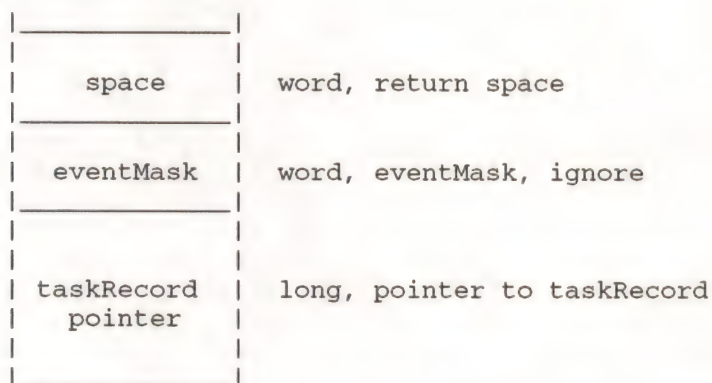


Figure 1-TaskMasterDA Stack Picture

The call returns a word value, the `taskCode`. The `space` and `eventMask` are self-explanatory. The book tells you that the `eventMask` is ignored, which makes sense since the host application has already gotten the event and you have already specified an `eventMask` in your NDA header, so you can use any value here. The `taskRecPointer` causes the confusion.

You do not pass a blank event record. When your NDA's action routine is

called, the Y and X registers contain a pointer to the current event record with which the NDA is working. TaskMasterDA is filling that taskRecord with some information, so you want to move it into your NDA's data area so you can work with it later:

```
phy
phx                               ; push the pointer that was passed to us
pushlong #NDArecord               ; the space in my NDA for the extended event record
pea      0
pea      16                       ; only 16 bytes, the original taskRecord size
_BlockMove
```

It is very important that you only move 16 bytes. TaskMasterDA can act erratically if the extended portion of the event record has been filled with nonsense values. This can happen if your NDA is running in an application that does not use the extended task record and you are copying non-task data into the extended portion of the task record. By the way, as you are debugging your NDA and you run into situations where the wmTaskData field values are weird, this is more than likely the problem.

Also remember to make sure the wmTaskMask field in your NDA's TaskRecord is set and the extended portions of the TaskRecord are zeroed out before your NDA starts running; you want to set all these fields in your NDA's INIT routine.

Now you can call TaskMasterDA:

```
pea      0                       ; return space
pea      $FFFF                   ; eventMask, ignored
pushlong #NDArecord             ; our NDA event record
_TaskMasterDA
pla                                       ; event code returned
```

You can then process the event in a convenient way. Remember again that TaskMaster has already done the control tracking for the controls in your NDA window. You have the same multiclick information, control handles and IDs.

Conclusion

TaskMaster is a wonderful thing that makes any programmer's job easier. So let it work for you. Learn the capabilities of the new fields and new controls, and take advantage of them. Let TaskMaster take care of the system details, while you concentrate on the features that make your application special.

Further Reference

-
- o Apple IIgs Toolbox Reference, Volumes 1 through 3
 - o Apple IIgs Technical Note #81, Extended Control Ecstasy
 - o Human Interface Note #8, Keyboard Equivalents

Apple IIgs
#85: Moving the Mouse

Written by: Matt Deatherage

July 1990

This Technical Note discusses moving the cursor on the screen without touching the mouse.

It is sometimes desirable to programmatically move the QuickDraw II cursor on the screen without requiring the user to touch the mouse. This can be effective, for example, in tutorial software that actually shows mouse actions such as pulling down menus and dragging windows.

There is not an easy or obvious way to do this in the toolbox. This is not a bad thing; it prevents overzealous programmers from zapping the mouse all over the screen for suspicious reasons. You must remember that the mouse belongs to the user, not to the application. If the user has put the mouse somewhere, it should only be a user's action that causes the cursor to move elsewhere. Most of the time that action is touching the mouse and physically moving it. Do not move the mouse except in response to a user-initiated command.

The most obvious way to move the mouse position--calling PosMouse with the new mouse position--is not sufficient; PosMouse does not update the current mouse position. When the mouse is next moved, a mouse interrupt comes through and the new deltas are added to the old mouse position, resulting in correct ReadMouse results after the mouse has been physically moved. Also, PosMouse does not update the cursor on the screen.

Faking Out the System

When you wish to move the mouse yourself, you are in effect replacing (or adding to) the standard mouse with a small programmatic mouse substitute--your code. This qualifies as a "device" and can be considered an Event Manager "device driver." You can then make the appropriate Event Manager call, FakeMouse. When calling FakeMouse, you supply all the mouse information yourself, allowing you to move the mouse, simulate button presses, and in general replace the mouse.

Further Reference

-
- o Apple IIgs Toolbox Reference, Volumes 1-3

Apple IIgs
#86: Risking Resourceful Code

Revised by: Matt Deatherage
Written by: C.K. Haun <TR>

March 1991
September 1990

This Technical Note covers considerations you need to keep in mind when using code resources.

Changes since September 1990: Now lists XCMD and XFCN resources as "Apple's code" and notes that other restrictions apply to them as well.

Code resources are wonderful things that can make your life better than it ever was before. Code resources are necessary when writing CDevs and can be very useful for control definition procedures, code modules for extensible programs like resource editors in fact, almost anywhere where you use regular compiled and linked code. But to do it right, you need to keep some rules in mind.

Apple's Code, Apple's Rules

The first code resources covered are the ones defined as fully supported by the System Software. These are rCtlDefProc (\$800C), rCodeResource (\$8017), rCDEVCode (\$8018), rXCMD (\$801E) and rXFCN (\$801F). Before looking at the specifics, this Note describes in general terms what happens when the Resource Manager loads a code resource.

When you call the Resource Manager with a request for a code resource (or when the system does, as with rCtlDefProcs), it loads it like a normal resource. The Resource Manager finds the resource in a resource map in the current search path, allocates a handle for the resource using the attributes in the resource attribute bits, and loads the resource into memory.

Now the Resource Manager examines the resConverter bit in the resource header. If this bit is set, indicating that this resource needs to be converted (as it should be for an rCtlDefProc), the Resource Manager checks its tables to see if a resource converter has been logged in (with the ResourceConverter call). For code resources, the correct converter has been logged in by the manager associated with that resource type. For example, the Control Manager logs in the code resource converter for rCtlDefProc resource type.

For code resources, InitialLoad2 is used to load the OMF from memory. Then the Resource Manager returns a handle containing a pointer to the start of the loaded, relocated code.

Rule 1: Code resources must be smaller than 64K

The code resource converter uses the InitialLoad2 function of the System Loader to load and convert code resources. That means that code resources are restricted in the same way that loading from memory is. One of these restrictions is that the code must be 64K or less.

Rule 2: Compiled and linked code only

Again, since InitialLoad2 is used to convert the code resource, the data must be in OMF format since InitialLoad2 expects to relocate standard load segments. When you prepare your code for inclusion in a code resource, compile and link the code as you normally would for a stand-alone program. Use the file produced by the linker for inclusion in your resource fork. You can use Rez to move the code from a data fork to your application's resource fork with a line in your resource description file similar to the following:

```
read rCodeResource (MyCodeIDNumber,locked,convert)"MyCompiledAndLinkedCode";
```

Rule 3: One segment please

Multiple segments are theoretically possible with code resources, but you have to manage memory IDs and the memory that the additional code segments use yourself. Since the code resource converter calls InitialLoad2, it uses the Memory Manager ID for the current resource application, and you cannot specify a different user ID directly. By changing your current resource application ID (by making an additional call to ResourceStartUp with a modified master ID, for example) you could manage multisegment code resources.

Rule 4: No dynamic segments

The dynamic segment mechanism does not work with code resources. Of course, your application can still use dynamic segments, but not code resource dynamic segments.

Rule 5: Set the right attributes

There are two sets of attributes you need to be concerned about for a code resource. The first set includes the standard resource attributes; the second set covers the attributes that the code itself has in the OMF image.

You need both sets to get the functionality you want. The resource attributes determine how the Resource Manager handles the resource. The OMF attributes control what InitialLoad2 does when it converts your code from OMF in a resource handle to relocated executable code.

Remember, you need to set both sets of attributes.

The resource attributes you need to set are locked and convert. The locked flag is necessary to prevent the resource from moving while InitialLoad2 processes it, and the convert flag is needed to signal the Resource Manager to call the code resource converter.

You must set the static OMF attribute, the others (like no special memory) you set as appropriate for your code in your application.

Rule 6: Know where to go

The handle you get back from the Resource Manager when you load and convert a code resource points to the beginning of the relocated and ready-to-execute code, not to the image of the code that is stored in the resource fork. So you can immediately jump to this code to execute it.

You can override this if you like clear the resConverter bit in the resource attributes. If this bit is zero, the Resource Manager does not call any resource converter (including the code resource converter).

Rule 7: Remember the Write

Keep in mind that any resource that uses a converter uses that converter both for reading and writing the resource. If you write out a code resource, the Resource Manager calls the Write routine for the code resource converter, which currently writes without doing any conversion it does not reconvert the code in memory back to OMF format. However, some converters (perhaps one you write) could reconvert the resource before writing it out.

Your Code, Your Rules

If you want to define your own code resource type (with a resource type of less than \$8000 and greater than 0) you may want to follow the same rules as the system code resources use. In fact, you can even use the same code resource converter, by using the ResourceConverter call with your resource type, and log the code resource converter as the converter to use with your resource type, like the following:

```
    pha
    pha                ; return space
    _GetCodeResConverter ; Misc Tools call to return the loader
*                       ; relocation code pointer
*                       ; (leave it on the stack for the next call)
    pea $0678          ; resource type you want to convert with this
*                       ; converter, any Application type you wish
    pea %01            ; add this converter to the Application
*                       ; converter list, and log this routine in
    _ResourceConverter
```

or you can do whatever you like with the resource, including not having a converter and doing all the relocation and memory management of the code yourself. This can give you the ability to add more functionality than the standard code resources provide dynamic segmentation is one feature you could implement if you want to handle all the details yourself.

Or, you can manage the code any way you want, but keep the built-in system functions in mind, and use as many of them as you can. Make your life simpler.

One Final Note

If one of your resources is marked convert and preload the Resource Manager only preloads that resource if the converter for that resource is logged in as a converter for that type. If the Resource Manager cannot find the converter, it does not preload the resource.

Further Reference

-
- o Apple IIgs Toolbox Reference, Volume 3
 - o GS/OS Reference
 - o HyperCard IIgs Script Language Guide
 - o HyperCard IIgs Technical Note #1, Corrections to the Script Language Guide
 - o Apple IIgs Sample Code #9, Lister

Apple IIgs
#87: Patching the Tool Dispatcher

Written by: Mike Lagae and Dave Lyons

September 1990

This Technical Note presents the Apple standard way to patch into the Apple IIgs Tool Dispatcher vectors.

This Note presents MPW IIgs assembly-language code which provides the Apple-standard way for utilities to patch and unpatch the Tool Dispatcher vectors. If all Tool Dispatcher patches follow this protocol, patches can be installed and removed in any order, without ever accidentally unpatching somebody who patched in after the one getting removed.

Using this protocol, each patch begins with a header in a standard form--a form recognizable by these routines (see PatchHeader). This way routines (like RemoveE10000) can scan through the list of patches and remove one from the middle.

If your patch is going to stay in the system until shutdown, use this standard patch protocol anyway. This way other utilities can still recognize your patch and scan past it to find the next one. This Note is not just to show you a way to patch the tool dispatcher--it's to show you the way. If you patch tool dispatcher vectors in any other way, you strip other utilities of their ability to remove their patches.

Of course, patching the Tool Dispatcher vectors slows down all toolbox calls, so you shouldn't patch the tool dispatcher without a pretty good reason. If you need to patch a toolbox function, it is usually better to do it by modifying a tool set's function pointer table instead of patching the dispatcher.

The code in this note is specific to the System tool dispatch vectors (\$E10000 and \$E10004), but the same technique is recommended for the User tool dispatch vectors--just change \$E10000 to \$E10008, \$E10004 to \$E1000C, and ToolPointerTable to UserToolPointerTable.

What Is This Stuff?

This Note presents the following four routines.

PatchHeader is the simplest patch function that obeys the protocol. This is where you put your own patch code.

InstallE10000 installs a patch into the patch chain. For example:

```
pushlong #PatchHeader
jsl InstallE10000
ply
ply ;remove the input parameter
bcc noError ;error in A
```

RemoveE10000 removes a patch from the patch chain. For example:

```
pushlong #PatchHeader
```

```

jsl RemoveE10000
ply
ply ;remove the input parameter
bcc noError ;error in A

```

CheckPatch determines whether the specified address is the starting address of a standard patch. For example:

```

pushlong #PatchHeader
jsl CheckPatch
ply
ply ;remove the input parameter
bcc validPatch

```

First, here are some comments and global equates.

```

*****
*
* Patch.e10000 - Routines to patch into the toolbox dispatch
*                vectors at $e10000 and $e10004.
*
* By Michael Lagae
* Software Quality Assurance
* GS Toolbox Test Team
*
* July 18, 1989
*
* Written for the MPW IIGS Assembler -- Version 1.1b1, 4/9/90
*
* Copyright 1989-1990 by Apple Computer, Inc.
*
*****

```

```

case yes
machine M65816
string asis
msb on
print on

```

```

export CheckPatch ; Check for a valid patch header.
export InstallE10000 ; Adds a patch into the toolbox
; vectors.
export RemoveE10000 ; Removes a toolbox dispatch vector
; patch.
export PatchHeader ; The simplest toolbox dispatch
; vector patch.

```

```

*****
* Equates - Various equates required by these routines.
*

```

```

versionNumber equ $0100 ; The version number of this library.

dispatch1 equ $e10000 ; The first toolbox dispatch vector.
dispatch2 equ $e10004 ; The second toolbox dispatch vector.
ToolPointerTable equ $e103c0 ; Pointer to the active System TPT.
UserToolPointerTable equ $e103c8 ; Pointer to the active User TPT.

; Error return values from the routines InstallE10000 and RemoveE10000

noError equ $0000 ; Value returned if no error occurs.
badHeaderError equ $8001 ; Patch header wasn't valid.
headerNotFoundError equ $8002 ; Header to remove wasn't in the

```

;

linked list.

PatchHeader is the standard shell for the actual patch code. Your code goes in here, at NewDispatch2. When you get control at NewDispatch2, the function number is in X and there are two RTL addresses on the stack (pushed after the function's parameters).

Your patch code does not care whether the tool call is being made through the \$E10000 or \$E10004 vector--in either case you get control with two RTL addresses on the stack.

* PatchHeader - Header required of all routines that will be patched
* into the toolbox dispatch vectors.
*

* Note: The code between next1Vector and NewDispatch2 must be included
* for all calls. The code below NewDispatch2 only needs to be
* included for patches that want to post patch the calls.
*

PatchHeader proc

entry next1Vector,next2Vector
entry dispatch1Vector,dispatch2Vector
entry NewDispatch1,NewDispatch2

next1Vector ; Where dispatch1 should go when
; finished.
; jml next1Vector ; (Filled in by InstallE10000).
next2Vector ; Where dispatch2 should go when
; finished.
; jml next2Vector ; (Filled in by InstallE10000).
dispatch1Vector ; Holds the JML instruction from
; \$e10000.
; jml dispatch1Vector ; (Filled in by InstallE10000).
dispatch2Vector ; Holds the JML instruction from
; \$e10004.
; jml dispatch2Vector ; (Filled in by InstallE10000).

anRtl rtl ; An RTL instruction. Its address
; will be
; pushed on the stack for dispatch1
; calls.

NewDispatch1 ; Entry point for dispatch1 toolbox
; vector.

phk ; Push program bank.
pea anRtl-1 ; Push the address of a RTL.

NewDispatch2 ; Entry point for dispatch2 toolbox
; vector.

; The following code should be included in the PatchHeader if the patch wants
; to perform post patching. This code will determine if the call that was
; made actually exists and if it does, post patching can occur. If the call
; doesn't exist, any pre-call routines can be executed, but the post patching
; shouldn't be attempted because the dispatcher will remove the second return
; address from the stack, thus not returning to your post patching routines.
; Stack equates for this routine.

aLong equ \$0001 ; A temporary long value.
oldDP equ aLong+4 ; Where the direct page is saved to.
oldTM equ oldDP+2 ; Where the tool call number is saved.

```

    phx                ; Save the call that's being made.
    phd                ; Save the current direct page.
    lda >ToolPointerTable+2 ; Get the TPT to determine the number
    pha                ; of tool sets loaded.
    lda >ToolPointerTable
    pha
    tsc                ; Set the direct page to the stack.
    tcd
    txa                ; See if this tool set exists.
    and #$00ff
    cmp [aLong]        ; Is it larger than the number of tool
;                        sets?
    bcs @noCall        ; JIF this tool set doesn't exist.
    asl a
    asl a
    tay                ; Now get the pointer to the FPT.
    lda [aLong],y
    tax
    iny
    iny
    lda [aLong],y
    sta aLong+2
    stx aLong
    lda oldTM          ; Get the function number.
    and #$ff00
    xba
    cmp [aLong]        ; Compare it to the number of entries in
;                        table.
; @noCall
    pla                ; Remove aLong from the stack.
    pla
    pld                ; Restore the original direct page.
    plx                ; Recover the tool number.

; At this point the carry flag is set if the tool call doesn't exist and clear
; if the tool call exists. No post patching must occur if the carry flag is
; set.

    jmp next2Vector    ; Go to the original $e10004 jump
;                        instruction.

    endp

```

* CheckPatch - Checks the passed toolbox dispatch vector to see if it
* points to a valid patch.

* Input: Passed via the stack following C conventions.
* newPatchAddr (long) - Address of the patch routine.

* Output:
* If newPatchAddr is a valid patch -
* Carry clear
* If newPatchAddr is not a valid patch -
* Carry set

CheckPatch proc

```

zprtl        equ $01                ; The address for the rtl on our direct
;                        page.
newPatchAddr equ zprtl+3            ; Address of patch (parameter to this

```

```

; routine).
; tsc ; Make the stack the direct page after
; saving
; phd ; the current direct page.
; tcd

; lda newPatchAddr+2 ; Simple check to check for a valid
; pointer.
; and #$ff00
; bne BadPatch ; Wasn't zero, can't be a valid pointer.

; lda [newPatchAddr] ; Check for the first JML instruction.
; and #$00ff
; cmp #$005c
; bne BadPatch

; ldy #$04 ; Check for the second JML instruction.
; lda [newPatchAddr],y
; and #$00ff
; cmp #$005c
; bne BadPatch

; ldy #$08 ; Check for the third JML instruction.
; lda [newPatchAddr],y
; and #$00ff
; cmp #$005c
; bne BadPatch

; ldy #$0c ; Check for the fourth JML instruction.
; lda [newPatchAddr],y
; and #$00ff
; cmp #$005c
; bne BadPatch

; ldy #$10 ; Check for the rtl and phk instructions.
; lda [newPatchAddr],y
; cmp #$4b6b
; bne BadPatch

; iny ; Check for the phk and pea instructions.
; lda [newPatchAddr],y
; cmp #$f44b
; bne BadPatch

; clc ; Calculate the address of the rtl
; instruction.
; lda newPatchAddr
; adc #$000f
; ldy #$13 ; Check for address of the rtl
; instruction.
; cmp [newPatchAddr],y
; bne BadPatch

GoodPatch
; pld ; Restore the direct page and report
; clc ; that it was a good patch.
; rtl

BadPatch
; pld ; Restore the direct page and report

```

```
sec
rtl
```

```
; that something was wrong.
```

```
endp
```

```
*****
* InstallE10000 - Sets the jump vector at $e10000 and $e10004 to point to
* the passed new toolbox dispatch vector patch. This routine
* also updates the linked lists so that more than one routine
* can be patched into the dispatch vectors.
*
* Input: Passed via the stack following C conventions.
* newPatchAddr (long) - Address of the patch routine.
*
```

```
* Output:
```

```
* If an error occurred -
* Carry set, Accumulator contains one of the following error codes:
* badHeaderError
* If no error occurred and patch was installed successfully -
* Carry clear, Accumulator contains zero.
*
```

```
InstallE10000 proc
```

```
oldPatchAddr equ $01 ; Address of existing patch.
zprt1 equ oldPatchAddr+4 ; The address for the rtl.
zpsize equ zprt1-oldPatchAddr ; Size of direct page we'll have on
; the stack.
newPatchAddr equ zprt1+3 ; Address of patch (parameter to
; this routine).
```

```
tsc ; Move the stack pointer to point beyond
sec ; the direct page variables that we'll
sbc #zpsize ; place on the stack.
tcs
phd ; Save the direct page register.
tcd ; Set the direct page.
php ; Disable interrupts
sei
```

```
pei newPatchAddr+2 ; Check if patch header is valid.
pei newPatchAddr
jsl CheckPatch
plx ; Remove the parameters from the stack.
plx
bcc @1 ; Report the badHeaderError if detected.
ldy #badHeaderError
jmp Exit
```

```
@1 lda >dispatch1 ; Set up the next1Vector in the new patch.
sta [newPatchAddr] ; The JML instruction and low byte.
lda >dispatch1+2
ldy #$02
sta [newPatchAddr],y ; The middle and upper bytes.

lda >dispatch2 ; Set up the next2Vector in the new patch.
ldy #$04
sta [newPatchAddr],y ; The JML instruction and low byte.
lda >dispatch2+2
ldy #$06
sta [newPatchAddr],y ; The middle and upper bytes.
```

```

;
    lda >dispatch1+3                ; See if there's already a patch in
;                                     dispatch1.
    and #$00ff
    sta oldPatchAddr+2
    pha                               ; High byte of possible header address.
    lda >dispatch1+1
    sec
    sbc #$0011
    sta oldPatchAddr
    pha                               ; Low byte of possible header address.
    jsr CheckPatch
    plx
    plx
    bcs First                        ; JIF this will be the first patch
;                                     installed.
;
    ldy #$08                          ; Set up the dispatch1Vector in the new
;                                     patch.
;
    lda [oldPatchAddr],y
    sta [newPatchAddr],y             ; The JML instruction and low byte.
    ldy #$0a
    lda [oldPatchAddr],y
    sta [newPatchAddr],y           ; The middle and upper bytes.
;
    ldy #$0c                          ; Set up the dispatch2Vector in the new
;                                     patch.
;
    lda [oldPatchAddr],y
    sta [newPatchAddr],y             ; The JML instruction and low byte.
    ldy #$0e
    lda [oldPatchAddr],y
    sta [newPatchAddr],y           ; The middle and upper bytes.
;
    bra PatchIt                      ; Now patch dispatch1 and dispatch2.
;
First ldy #$08                        ; Set up the dispatch1Vector in the new
;                                     patch.
    lda >dispatch1
    sta [newPatchAddr],y           ; The JML instruction and low byte.
    ldy #$0a
    lda >dispatch1+2
    sta [newPatchAddr],y           ; The middle and upper bytes.
;
    ldy #$0c                          ; Set up the dispatch2Vector in the new
;                                     patch.
;
    lda >dispatch2
    sta [newPatchAddr],y           ; The JML instruction and low byte.
    ldy #$0e
    lda >dispatch2+2
    sta [newPatchAddr],y           ; The middle and upper bytes.
;
PatchIt
    clc                               ; Calculate the address of the new
;                                     dispatch2.
    lda newPatchAddr
    adc #$0015
    sta newPatchAddr
    xba
    and #$ff00                       ; Mask in the JML instruction.
    ora #$005c
    sta >dispatch2                  ; The JML instruction and low byte.
    lda newPatchAddr+1
    sta >dispatch2+2                ; The middle and upper bytes.

```

```

;
; Calculate the address of the new
; dispatch1.
        lda newPatchAddr
        sbc #$0004
        sta newPatchAddr
        xba
        and #$ff00           ; Mask in the JML instruction.
        ora #$005c
        sta >dispatch1      ; The JML instruction and low byte.
        lda newPatchAddr+1
        sta >dispatch1+2    ; The middle and upper bytes.

        ldy #noError        ; Report that all went well.

Exit    plp                 ; Restore the interrupt state.
        pld                 ; Restore the previous direct page
;                                     register.
        tsc                 ; Restore the stack pointer.
        clc
        adc #zpsize
        tcs
        tya                 ; Value to return.
        beq @noerr          ; Report that there was an error.
        sec
        rtl
@noerr  clc                 ; Report that there was no error.
        rtl

        endp

```

```

*****
* RemoveE10000 - Removes the specified patch from the dispatch1 and dispatch2
*               vectors and updates the linked lists for the remaining
*               toolbox patches.
*
* Input: Passed via the stack following C conventions.
*        patchToRemove (long) - Address of the patch to remove.
*
* Output:
*        If an error occurred -
*            Carry set, Accumulator contains one of the following error codes:
*                badHeaderError
*                headerNotFoundError
*        If no error occurred and patch was removed successfully -
*            Carry clear, Accumulator contains zero.
*

```

RemoveE10000 proc

```

patchDispAddr    equ $01           ; Address of existing patch (and 1
;                                     extra byte).
prevHeader       equ patchDispAddr+5 ; Used to search through the
;                                     linked list.
zprtl            equ prevHeader+4   ; The address for the rtl.
zpsize          equ zprtl-patchDispAddr ; Size of direct page we'll have
;                                     on the stack.
patchToRemove    equ zprtl+3       ; Address of patch (parameter to
;                                     this routine).

        tsc                 ; Move the stack pointer to point beyond
        sec                 ; the direct page variables that we'll

```

```

    sbc #zpsize                ; place on the stack.
    tcs
    phd                        ; Save the direct page register.
    tcd                        ; Set the direct page.
    php                        ; Disable interrupts
    sei

    pei patchToRemove+2        ; Check if patch header we were asked to
    pei patchToRemove          ; remove is a valid header.
    jsl CheckPatch
    plx                        ; Remove the parameters from the stack.
    plx
    bcc @1                    ; Report the badHeaderError if detected.
    ldy #badHeaderError
    jmp Exit

@1    clc                    ; Create the JML instruction that would
    ;                          exist
    lda patchToRemove          ; if the patchToRemove was installed.
    adc #$0011
    sta patchDispAddr+1
    lda patchToRemove+2
    sta patchDispAddr+3
    lda patchDispAddr          ; Mask in the JML instruction.
    and #$ff00
    ora #$005c
    sta patchDispAddr

    cmp >dispatch1            ; Check if the patch to remove is the
    ;                          first
    bne NotFirstOne          ; patch installed.
    lda >dispatch1+2
    cmp patchDispAddr+2
    bne NotFirstOne

    lda [patchToRemove]        ; Restore the Dispatch1 vector.
    sta >dispatch1
    ldy #$02
    lda [patchToRemove],y
    sta >dispatch1+2

    ldy #$04                  ; Restore the Dispatch2 vector.
    lda [patchToRemove],y
    sta >dispatch2
    ldy #$06
    lda [patchToRemove],y
    sta >dispatch2+2

    bra NoErr                ; Everything went well.

NotFirstOne
    sec                    ; Assume that whatever is in dispatch1 is
    lda >dispatch1+1        ; patch and get the address of its header.
    sbc #$0011
    sta prevHeader          ; Low and middle bytes.
    lda >dispatch1+3
    and #$00ff
    sta prevHeader+2        ; Upper byte of header address.

@loop pei prevHeader+2        ; Check if it really is a valid header.
    pei prevHeader

```

```

jsl CheckPatch
plx                               ; Remove the parameters from the stack.
plx
bcc @2                             ; Report that the patch that we asked to
ldy #headerNotFoundError         ; remove wasn't found.
bra Exit

@2    lda [prevHeader]             ; See if this patch points to patch we
      cmp patchDispAddr           ; want to remove.
      bne @nope
      ldy #$02
      lda [prevHeader],y
      cmp patchDispAddr+2
      bne @nope

      lda [patchToRemove]         ; Restore the next1Vector.
      sta [prevHeader]
      ldy #$02
      lda [patchToRemove],y
      sta [prevHeader],y

      ldy #$04                     ; Restore the next2Vector.
      lda [patchToRemove],y
      sta [prevHeader],y
      ldy #$06
      lda [patchToRemove],y
      sta [prevHeader],y

      bra NoErr                   ; Everything went well.

@nope ldy #$02                     ; Get the address of the next patch
;                                         header.
      lda [prevHeader],y
      tax
      lda [prevHeader]
      sta prevHeader
      stx prevHeader+2

      sec
      lda prevHeader+1
      sbc #$11
      sta prevHeader
      lda prevHeader+3
      and #$00ff
      sta prevHeader+2

      bra @loop                   ; Now check this header.

NoErr ldy #noError                 ; Report that all went well.

Exit  plp                           ; Restore the interrupt state.
      pld                           ; Restore the previous direct page
;                                     register.
      tsc                           ; Restore the stack pointer.
      clc
      adc #zpsize
      tcs
      tya                           ; Value to return.
      beq @noerr
      sec                           ; Report that there was an error.
      rtl

```

```
@noerr clc  
      rtl
```

```
; Report that there was no error.
```

```
      endp
```

```
end
```

Further Reference

- o Apple IIgs Toolbox Reference
- o Apple IIgs Technical Note #73, Using User Tool Sets

Apple IIgs

#88: The Page One Stack in a 16-Bit World

Written by: Dave Lyons

September 1990

This Technical Note clarifies the protocol for moving the stack pointer in and out of page one.

On page 13 of the Apple IIgs Firmware Reference, under "Save the value of the native-mode stack pointer," there is a code sample showing how to switch to the page-one stack by setting the stack pointer to \$01xx, where xx is the contents of EMULSTACK at \$01/0100.

However, the manual does not warn you about moving the stack pointer from page one to another area. When you do that, you must store the low byte of the stack pointer at EMULSTACK before moving the stack pointer out of page one. If you do not save the page-one stack properly, interrupt routines or some toolbox calls may destroy a part of the page one stack that you go back to later, expecting that return addresses are still there.

Note: If the auxiliary-memory stack and zero page are in use, you must use \$01/0101 instead of \$01/0100. See the Apple IIe Technical Reference Manual, pp. 153-154

Further Reference

- o Apple IIgs Firmware Reference
- o Apple IIe Technical Reference Manual

Apple IIgs

#89: MessageByName--Catchy Messages

Written by: Dan Strnad & Dave Lyons

September 1990

This note clarifies MessageByName and provides examples of creating and retrieving a named message.

Did You Say You Want To Get A Message?

All you have to do is ask. Apple IIgs Toolbox Reference, Volume 3 already tells you how. Here's what the fine print says: with the createItFlag set to FALSE and the name of the message you are after in the nameString, you call MessageByName. What's unclear in the manual is that if the message was found, no error is returned, the createFlag is returned as FALSE, and messageID contains the ID you can pass to MessageCenter to retrieve the contents of the message. Here's an example of MessageByName in use.

The following code creates a named message.

```
CreateNamedMessage
    pha
    pha
    pea 1                ;create it
    pushlong #MsgBlock
    _MessageByName      ;function $1701
    pla
    sta myMsgID         ;keep the ID if you want
    pla                 ;check the createFlag if
;                       you want
    ...

MsgBlock    dc.w MsgBlockEnd-MsgBlock
            dc.b 28, 'XYZ Software:My Cool Product' ;Pascal-style string
            ... more data goes here

MsgBlockEnd
```

The following code retrieves the message.

```
    pha
    pha
    pea 0                ;don't create message
    pushlong #MsgBlock
    _MessageByName      ;function $1701
    ply                 ;keep id of existing
;                       message
    pla                 ;createFlag (ignore)
;                       ;carry set if an error
;                       occurred

    pea 2                ;MessageCenter action:GET
    phy                 ;message ID for
;                       MessageCenter, below
    pha
```

```

;
    pha                                ;space for NewHandle
    lda #0                             ;result
    pha                                ;size of handle (0)
    pha
    lda MyID                           ;ID for empty
    pha
    pha                                ;handle attributes (0)
    pha                                ;no special location
    _NewHandle
    lda 3,s
    sta mcHandle+2
    lda 1,s
    sta mcHandle
;
    _MessageCenter
;

    lda mcHandle+2
    pha
    lda mcHandle
    pha
    phd
    tsc
    tcd
    ldy #2
    lda [3],y
    tax
    lda [3]
    sta 3
    stx 5

```

* now read data from the message at [3]

```

    ldy #$xxxx                          ;index past the name
;                                         string
    lda [3],y
    ...
    pld
    pla
    pla

    lda mcHandle+2
    pha
    lda mcHandle
    pha
    _DisposeHandle

```

noMessage ...

mcHandle dc.l 0
myMsgID dc.w 0

MessageByName is available in Tool Locator versions 3.0 and later (System Software 5.0 and later).

Further Reference

o Apple IIgs Toolbox Reference, Volumes 2-3

Apple IIgs
#90: 65816 Tips and Pitfalls

Revised by: Matt "Matt" Deatherage
Written by: Dave "Dave" Lyons

March 1991
September 1990

This Technical Note presents short 65816 assembly language examples illustrating pitfalls and clever techniques.

Changes since November 1990: Added more explanations about the JSL table and corrected a comment.

Dispatching Through an Address Table

The 65816 has a JSR (\$aaaa,X) instruction for calling a selected subroutine from a table of addresses, but it has no JSL (\$aaaa,X) instruction. If you need to dispatch to one of several routines that are not all in the same bank,

you need an approach like the following. The idea is to perform a JSL to a routine which does a long jump by pushing a three-byte "RTL address" on the stack and then doing an RTL.

```
        jsl LngJump          ;go jump to the routine
        ...

LngJump  asl a                ;take routine number in A and
        asl a                ; multiply it by 4
        tax                 ;put table index into X
        lda table+1,x       ;get "middle" word of address
        pha                 ; and push it
        lda table,x         ;get low word and
        dec a               ; decrement it by one
        phb                 ;push a single throw-away byte
        sta 1,s             ;store over low two of the 3 bytes
        rtl                 ;transfer control to the routine
table    dc.l routine1      ;table of 4-byte subroutine addresses
        dc.l routine2
        dc.l routine3
        ...
```

This code is correct because RTL pulls three bytes off the stack and increments the two low bytes without incrementing the high byte.

Note: This approach to a table-based JSL is more flexible than JML (\$XXXX) because it does not require any fixed-location storage or bank zero space, other than the stack.

On the other hand, the following code is not correct. The approach here is to make a table of addresses minus one.

```
        asl a                W
        asl a                R ;multiply index by 4
```

```

tax          O ; and put it in X
lda table+1,x N ;get the "middle" word
pha         G ; and push it
lda table,x  ! ;get the low word
phb        W ;push a single throw-away byte
sta 1,s     R ;store over low two bytes
rtl        O ;transfer control to the routine
table      dc.l routine1-1 N ;table of 4-byte addresses minus one
           dc.l routine2-1 G
           dc.l routine3-1 !
           ...

```

This second sample code fragment fails if any of the routines in the table comes at the first byte of a bank. For example, if routine1 is at \$060000, the address pushed is \$05FFFF, and RTL transfers control to \$050000, not \$060000.

Dereferencing Handles Without Direct Page Space

When your code gets called with the D register undefined, you must not use direct page addressing without setting D to a known good value. Preserving and restoring locations on the caller's direct page is not reliable, because D could be pointing at bytes below the stack pointer (which can be destroyed by interrupts) or even at the \$C0xx soft switches (that would make your direct page accesses accidentally fiddle with hardware).

A common way to get temporary direct page space is to point D at part of your stack. This following code dereferences a handle stored in the A and X registers (if the handle is \$E01234 and refers to a block of memory at \$056789, then on entry A=\$00E0 and X=\$1234, and on exit A=\$0005 and X=\$6789).

```

phd          ;save caller's direct-page register
pha         ;push high word of handle
phx        ;push low word of handle
tsc        ;get stack pointer in A
tcd        ;and put it in D
lda [1]     ;get low word of master pointer (no ",Y"! )
tax        ; and put it in X
ldy #$0002 ;offset to high word of master pointer
lda [1],y   ;get high word
ply        ;remove low word of handle
ply        ; and high word
pld        ;restore the caller's direct-page register

```

Direct page addressing isn't the only way to address through pointers. Here's the same routine as before, but using the Data Bank register (B) instead of fiddling with D. (Note that handles do not have to be in bank \$E0 or \$E1, although they usually are.)

```

phb          ;save caller's data bank register
pha         ;push high word of handle on stack
plb        ;sets B to the bank byte of the pointer
lda |$0002,x ;load the high word of the master pointer
pha         ; and save it on the stack
lda |$0000,x ;load the low word of the master pointer
tax        ;and return it in X
pla        ;restore the high word in A
plb        ;pull the handle's high word high byte off the
           ; stack
plb        ;restore the caller's data bank register

```

You don't have to switch into Native mode just to do an eight-bit operation with long addressing. Most 65816-specific instructions and addressing modes work in emulation mode in approximately the same way they work in eight-bit native mode. See the "Further Reference" for details.

Further Reference

- o Apple IIgs Hardware Reference
- o Programming the 65816, Including the 6502, 65C02 and 65802 (Eyes and Lichty, 1986, Brady)

Apple IIgs

#91: The Wonderful World of Universal Access

Revised by: Dave Lyons

May 1992

Written by: Don J. Brady, Matt Deatherage, & Ron Lichty

September 1990

This Technical Note discusses how your applications can be compatible with Universal Access software.

CHANGES SINCE JULY 1991: Added caution against reading the keyboard with interrupts disabled.

WHAT'S "UNIVERSAL ACCESS?"

Universal Access is the name given to software components designed to make Apple computers (in this case, the Apple IIgs) more accessible to people who might have difficulty using them. The Apple IIgs is very dependent on graphic objects, a keyboard and mouse; not all people can use these things very easily.

There are several components to Apple's Universal Access software:

- o CloseView. CloseView magnifies the Apple IIgs screen so that it's more easily seen by those with visual impairments. The hardware screen contains a magnification from two to twelve times as large as the "real" 32K Super Hi-Res graphics screen.
- o Video Keyboard. Video Keyboard is a New Desk Accessory that emulates a keyboard. A picture of a keyboard appears on the screen; a mouse-down event in any "key" makes Video Keyboard post a key-down event, so you can use a pointing device as a keyboard. ADB hardware is available to allow people to use head gear or other devices instead of mice; Video Keyboard lets these same devices replace the keyboard as well.
- o Easy Access. Easy Access comes in two parts: Sticky Keys and Mouse Keys. Sticky Keys makes the keyboard easier to use for those who have trouble pressing more than one key at a time; while Sticky Keys is activated, modifier keys may be released and still apply to the next keystroke. Mouse Keys allows the numeric keypad to be used as a mouse substitute. Sticky Keys and Mouse Keys are included in all ROM 03 Apple IIgs computers. The software versions allow all Apple IIgs computers to provide these functions, and provide additional icon feedback (in the upper right menu bar) for Sticky Keys.

HOW IT WORKS (ACCESS NOTHING AND CHECKS FOR FREE)

Universal Access generally works by replacing Apple IIgs toolbox functions. For example, CloseView patches QuickDraw so you do not draw to the actual screen, but to another buffer that CloseView can then magnify. Video Keyboard patches the Window Manager so that its keyboard window is always frontmost and fully visible (and accessible). Easy Access uses the ADB tools and the Event Manager to alter the way the hardware responds.

Since Universal Access changes the way the tools behave, your applications do

not have to work very hard to be accessible to a broad range of physically challenged people. Just by following the rules, you have an accessible application. There are, however, a few guidelines you should keep in mind when designing your programs to make them as accessible as they can be.

UNIVERSAL ACCESS COMPATIBILITY GUIDELINES

- o Don't disable interrupts and then try to read the keyboard. Easy Access on ROM 1 works at the Apple Desktop Bus level--if ADB interrupts are not being serviced, no keypresses will show up at \$C000/\$C025. Even Reset will not work, so the user may have to power down to regain control of the machine.
- o Try to avoid using modal dialogs. Not only do lots of modal dialogs make for a cumbersome interface for everyone, they are especially annoying to those who have to move the mouse to a lot of OK buttons. More importantly, users cannot open NDAs like Video Keyboard while modal dialogs are frontmost.

Video Keyboard can also be dragged in front of modal dialogs. If you are in the habit of using QuickDraw calls to draw items in Dialog Manager modal dialogs instead of creating custom dialog userItems, Video Keyboard users can drag the keyboard window in front of your dialog and erase the items (since the only items redrawn are those redrawn by the Dialog Manager's update routine). You can easily test this in all of your dialogs by obscuring each dialog with the Video Keyboard window a piece at a time, then moving Video Keyboard away, to be sure that all areas are completely redrawn.

Let's say, for example, that you have a custom text item that changes between invocations of the same modal dialog. You might choose to draw the text yourself with LETextBox2 after creating the dialog with GetNewModalDialog but before letting the Dialog Manager handle events with ModalDialog:

```
phx                ; port: hi word from GetNewModalDialog
pha                ; port: lo word from GetNewModalDialog
_SetPort

lda OurText+2     ; pointer to text to draw in modal dialog
pha
lda OurText
pha
lda OurTextLength ; Text length
pha
pea OurTextRect>>16 ; Text rectangle
pea OurTextRect
pea 0002          ; Text justification (2 = fill)
_LETextBox2
```

To be Universal Access-friendly, you would, instead, implement a userItem routine like the following:

```
; . . . . .
DrawDialogText
;
; DrawDialogText draws text pointed to by OurText into the Dialog.
; This userItem routine is called only by the Dialog Manager,
; when it's implementing/updating the dialog.
; . . . . .
```

```

lda >OurText+2          ; pointer to text to draw in modal dialog
pha
lda >OurText            ; (long addressing: data bank unknown)
pha
lda >OurTextLength     ; Text length
pha
pea OurTextRect>>16    ; Text rectangle
pea OurTextRect
pea 0002                ; Text justification (2 = fill)
_LETextBox2

lda 1,s                ; get return address
sta 7,s                ; move to proper location
lda 2,s                ; above input parameters
sta 8,s

pla                    ; move stack pointer up
pla                    ; to new return address location
pla
rtl

```

It will be called as a result of adding a template item like the following to the dialog template (note use of Item Value for the text length, since template Value fields are not used by userItems):

```

TextTemplate  dc.w 3                ; ID
OurTextRect   dc.w TTop, TLeft, TBottom, TRight
               dc.w UserItem+ItemDisable ; Type
               dc.l DrawDialogText    ; Pointer to our userItem
               ; routine
OurTextLength ds.w 1                ; Text length (cheap place
               ; to put it)
               dc.w 0000              ; Item flag
               dc.l 00000000          ; Item color

```

Note that this is a simple example of a custom item routine; if you really had custom text that changed from invocation to invocation, you could use the existing Dialog Manager ParamText and longStatText2 item mechanisms.

- o Use the Event Manager routines for event information. Do not access any hardware directly or use the lower-level Miscellaneous Tools routines for user event information--you steal that information from Universal Access. For example, use the Event Manager routine GetMouse to find the mouse location. Do not use ReadMouse or you steal mouse movement information from Universal Access.
- o Call GetNextEvent or TaskMaster often. Long delays between calls do not let NDAs like Video Keyboard get events. If you cannot make these calls, at least call SystemTask.
- o Do not assume that the hardware location of the screen is \$E12000. Universal Access components that manipulate the entire screen (like CloseView) move the virtual screen so the hardware can be used for the magnified screen image.

To find the screen location, look at the ptrToPixImage field in a grafPort after calling OpenPort (or in your window's window record after NewWindow). The image pointer gives the correct location of the screen.

Assuming the current port is on screen, the following code finds the

ptrToPixImage value:

```
pha
pha                ;made space for port pointer
_GetPort
phd                ;save direct page location
tsc
tcd                ;port pointer is now at 3..6 on direct page
ldy #4             ;offset to high word of ptrToPixImage
lda [3],y          ;got high word
tax                ; in X
ldy #2             ;offset to low word of ptrToPixImage
lda [3],y          ;got low word
tay                ; in Y
pld                ;restored direct page location
pla
pla                ;removed port pointer
```

The X and Y registers now contain the base address of the screen.

- o Do not assume things about being the frontmost window. Even if FrontWindow says you have the frontmost window, your visRgn may have pieces missing. For example, the title bar of your window may be partially under the menu bar. Or there may be a floating "windoid" (like Video Keyboard's window) over part of your window.

For these reasons you should not draw directly to the screen without first examining your window's visRgn. Do not just check for rectangularity--your visRgn could be rectangular and parts of your window still be obscured. If you use QuickDraw for all your drawing, QuickDraw automatically clips drawing activity to be entirely within the visRgn, so this is not a problem.

- o Don't access QuickDraw data directly; use QuickDraw routines instead. For example, to access SCB data, use the QuickDraw routines GetSCB and SetSCB instead of reading the hardware at \$E19D00. CloseView may have those SCBs changed to reflect a magnified portion of the screen. Also use GetColorEntry, SetColorEntry, GetColorTable, and SetColorTable. Don't access the hardware directly.
- o Try to allocate memory after starting the tools. If you want to allocate memory before starting tools, do not use special memory. (Set the attrNoSpec bit in the attributes.)

Further Reference

- o Apple IIgs Toolbox Reference
- o Apple IIgs Firmware Reference
- o Apple II Video Overlay Card Development Kit (APDA)

Apple IIgs

#92: Twisted Tales of TextEdit

Revised by: Dave Lyons

December 1991

Written by: C.K. Haun <TR> and Dave Lyons

September 1990

This Technical Note discusses some undocumented features and some bugs in the TextEdit tool set through System Software 5.0.4.

Changes since November 1990: Noted that a non-control TENew creates a Text Edit record for the current port.

TENew

TextEdit records you create with TENew are always tied to the current port at the time of the TENew call, whether or not the fNotControl bit is set. (For TextEdit controls, NewControl2 is the preferred call.)

TEInsert

Using the TEInsert call on an invisible TextEdit record causes the screen to scroll, exactly as if the TextEdit record were visible.

If you use LETextBox2 style text as input for a TEInsert call, any style change information contained at the end of the LETextBox2 text is ignored. To ensure that the style change is not ignored, append an additional character at the end of the block, then delete (with TEsSetSelect and TEdDelete) the extra character after the TEInsert call.

TEGetText

The documentation for TEGetText says that a dataFormat value of \$4 returns the text as "Formatted for input to LineEdit LETextBox2". This is not a reliable return method-this call may or may not succeed. Greater chance for success occurs with less than 4,000 characters in the TextEdit record.

TEGetText also supports getting just the text of the current selection range. Adding \$0020 (onlyGetSelection) to the number passed in bufferDescriptor returns the text of the current selection. This technique does not work with data format LETextBox2, but does work with all other formats. Also, there is no corresponding bit for the associated style record, so you cannot get the style for just the current selection this way, if you request style information you get a styleRef for the entire TextEdit record.

TEClick

Using TEClick or TestControl on an inactive record currently causes that record to activate.

TERuler

Pixel tabbing values must all be greater than zero or TextEdit loops infinitely on a tab.

TEGetRuler & TEsSetRuler

TERuler, for the default ruler or any ruler that uses a tabType value of \$1

returns a ruler four bytes longer than described in the documentation. The extra four bytes are all \$FF, and they are the terminator characters for tabType \$2 rulers. Expand your buffers by four bytes to prevent overwriting any data. TextEdit also expects the additional information on a TEsSetRuler call, so you should pad your ruler with four \$FF bytes if you are using a type \$1 ruler.

TEsSetText

Passing a zero-length class one input string (a word length string with the word set to zero) to TEsSetText causes TextEdit to crash.

TEPaintText

TEPaintText currently prints colored text in only four colors.

It's Not Dirty, It's Text

There has been some confusion about determining if a TextEdit record has been changed. The documentation has been a little vague, and the process itself has misled some people. Here is The Truth: there is a TextEdit dirty flag, and you can use it and rely on it to tell you when a TextEdit record has changed.

The TextEdit dirty flag is bit 6 (fRecordDirty in the E16.TextEdit interface file) of the ctlFlag byte. This has caused some confusion because the ctlFlag byte is at offset \$12 in the control definition template, and it is at offset \$10 in the TextEdit or Control record. Just remember that it is not in the same place in the record as it is in the template.

If it is set, then the TextEdit or Control record has been changed since the last time the dirty bit was cleared. The dirty bit is clear initially when you create the TextEdit or Control record. Anytime after that, if the user enters text into the TextEdit record, TextEdit sets the dirty flag. It is up to your application to clear the dirty flag; TextEdit has no way of knowing when you've saved or cleared data.

Further Reference

-
- o Apple IIgs Toolbox Reference, Volume 3

Apple IIgs

#93: Compatible Printing

Revised by: Matt Deatherage May 1992
Written by: Matt Deatherage September 1990

This Technical Note discusses printing on the Apple IIgs and how you can make your printing code more compatible.

CHANGES SINCE SEPTEMBER 1990: Added a note about expecting print records to keep the same attributes across Print Manager calls. Added the StyleWriter's iDev value.

HOW DOES PRINTING WORK ANYWAY?

There are, in general, two types of printing done on the Apple IIgs. The first kind is "desktop" printing, which uses the Apple IIgs Print Manager to render images created by QuickDraw II onto an output device. The other kind of printing is "text" printing, which is similar to the way classic Apple II applications print--you send ASCII text somewhere and a printer prints it as ASCII text. This printing normally involves no graphics and is very quick.

This Note covers both types of printing, and by understanding the internals and the methods used to print, you can avoid compatibility headaches in the future.

DESKTOP PRINTING

Desktop printing uses the Apple IIgs Print Manager. The process is described in detail in the Print Manager chapter of the Apple IIgs Toolbox Reference, and usually consists of a simple print loop:

```
Open a document (PrOpenDoc)
  Open a page (PrOpenPage)
  Draw or Image the page in your favorite way
  Close the page (PrClosePage)
  Repeat for each page
Close the document (PrCloseDoc)
Print the document if it's spooled (PrPicFile)
```

Note that you should ALWAYS call PrPicFile at the end of your print loop. It completes the printing process, even for immediate or draft printing.

There's one real secret about the Print Manager that can cloud your understanding of printing--the Print Manager doesn't actually do anything. It loads, unloads, and keeps track of printer drivers and port drivers and performs some necessary housekeeping, but that's about it. Many people believe that the Print Manager is responsible for all imaging, managing documents, managing a printing grafPort and such, but it's not. (The myth is perpetuated by the Toolbox Reference which refers to these functions as handled by the Print Manager.) In fact, these functions are handled by printer drivers.

You actually call the printer driver for all of the routines in the print loop; all the Print Manager does is make sure the driver is loaded and

dispatch to it. Therefore, most of the compatibility issues you have with printing are not with the Print Manager, but with printer drivers.

DEALING WITH THE PRINT RECORD

It's the printer driver's job to get information about a printing job from the user (it's the printer driver that handles the style and job dialog boxes, since the Print Manager cannot generically know what style and job options any printer can support), keep track of it, and print the document using those settings. Those settings are kept in a data structure associated with a document known as a print record.

Apple had only released two printer drivers at the time the first volume of the Toolbox Reference was published, and therefore the descriptions of the print record in that volume tend to be absolute. For example, the `iDev` field is documented as "one for an ImageWriter and three for a LaserWriter." In fact, the `iDev` field is the only method of print record interpretation available and there are several values for it:

```
$0001 = ImageWriter
$0002 = ImageWriter LQ
$0003 = LaserWriter
$0004 = Epson
$0065 = StyleWriter
$8001 = Generic dot-matrix (interprets the style subrecord like the
      ImageWriter driver)
$8003 = Generic laser printer (interprets the style subrecord like the
      LaserWriter driver)
```

If you have checks in your code like "If it's not \$0001, it must be a LaserWriter," you have problems with most of the other printer types.

The \$8000 and greater `iDev` values are defined for third-party printer drivers. The printer driver has no way other than the print record to keep track of values for a given print job, so it has to store all such information in the print record. If all third-party drivers use proprietary style subrecord formats, no applications can read or set any of those values. Those drivers which can use the compatible \$8000 and greater `iDev` values indicate to applications that the definitions in Toolbox Reference for the ImageWriter and LaserWriter drivers apply to these drivers as well. `iDev` values of \$0002 or \$0004 also interpret the style subrecord as the ImageWriter driver does.

PRINT RECORD RULES

Remember: the print record is the only way the printer driver has to maintain information about a particular job. The print record belongs to the user, the document, and the printer driver--NOT the application. Here are some rules for staying out of print record trouble.

- o Always call `PrValidate` when changing fields in the print record. Even if a driver interprets the style subrecord like the ImageWriter driver, it may not support all the ImageWriter's style features (e.g., color printing). Calling `PrValidate` every time you change something in the print record gives the printer driver a chance to look at the havoc you've wreaked and correct it if necessary.

You do not always get a feature you want. If a printer does not support color printing, you can set the "color" bit all day long and `PrValidate` clears it every time. You should be prepared for a new printer driver that does not support the features you want, and inform the user that the feature is not supported by this printer.

- o Do not patch PrValidate to make it ignore bogus values in the print record unless instructed to do so by the printer driver author.
- o Never, never tread on reserved fields in the print record. If you find a particular driver storing useful values some place, forget it. This is the only place a driver has to store information about a print job and some of it is not going to be supported.

In particular, never try to interpret any values you may find in the printX subrecord of the print record. This subrecord is for the private use of printer drivers. Although printX is currently the worst compatibility risk, you must not tamper with other reserved fields.

- o Don't assume that the print record will keep the same memory attributes across calls to the Print Manager (and therefore the printer driver). Specifically, don't assume that a print record will stay locked across calls to the Print Manager.
- o If you want to learn more about printing, learn how printer drivers work. The specifications are in Apple IIgs Technical Note #35, appropriately entitled "Printer Driver Specifications." An understanding of how printer drivers do their work is an understanding of how printing works.

TEXT PRINTING

Text printing generally uses the built-in ASCII mode of most dot-matrix printers to print text quickly and efficiently.

Desktop printer drivers often have a "draft" mode, where they print text immediately instead of imaging it in the appropriate font and style. This is accomplished by intercepting low-level QuickDraw II routines called bottleneck procedures. When QuickDraw is called to draw text, the printer driver gets control instead and sends the text to the printer.

Although this is useful to users of desktop printer drivers, it is not a required feature of any printer driver, and those that do implement it each do so in their individual way. For example, the LaserWriter driver doesn't support this model of "draft" printing because the LaserWriter is normally a PostScript(R) device--sending straight ASCII to it doesn't necessarily work.

To imitate the way classic Apple II applications print, your application prompts the user for some device through which to print, and ASCII characters are sent through that device. There are a few ways to do this.

USING THE PRINT MANAGER

You can still use the Print Manager to print in ASCII mode by bypassing the printer driver. Simply use the Port Driver to send ASCII characters to the given target device with the PrDevWrite call. The specifications for Port Driver calls are in Apple IIgs Technical Note #36, also appropriately entitled "Port Driver Specifications." You make port driver calls as if they were Print Manager calls.

Although this method has been used, Apple does not recommend it. If the selected port driver is a network driver, this method is troublesome.

USING THE TEXT TOOLS

By using the Apple IIgs Text Tools, you can ask the user what slot to print

through and send ASCII characters to that slot or port. Although this is better than using the Port Driver, it still has problems. The Text Tools cannot be fully GS/OS Slot Arbiter compatible; therefore, there might be GS/OS devices accessible to the user to which your application does not let him print. Also, it's difficult to detect which slots really have Text Tools' devices without knowing about Apple II firmware, and prompting the user for a slot number invites trying to print to the disk firmware, which usually justs reboot the machine (unceremoniously).

USING GS/OS

GS/OS supports character drivers, such as printer interfaces, and using them is the best way to handle ASCII printing. GS/OS supports loaded drivers for character devices if you have them, and generates drivers for character devices it can recognize. In addition, GS/OS drivers have identification words so you can prompt with real messages instead of cryptic slot numbers.

You can use the GS/OS call DInfo to loop through all drivers and prepare a list of character drivers. You can then change their device IDs into text phrases, place them in a list, and prompt the user to select one. This call usually results in a list such as "Printer port, Modem port, Remote Print Manager, Printer interface, Text screen [the Console driver]." You may wish to change the names of the devices slightly to make the choice easier (e.g., "network printer" instead of "Remote Print Manager").

Apple strongly recommends using GS/OS for ASCII printing from 16-bit applications.

NOTE : The Remote Print Manager (.RPM) device driver in System Software 5.0 to 5.0.2 has a bug which causes character loss. System Software 5.0.3 fixes this bug.

Further Reference

- o Apple IIgs Toolbox Reference
- o GS/OS Reference
- o Apple IIgs Technical Note #34, Low-level QuickDraw II Routines
- o Apple IIgs Technical Note #35, Printer Driver Specifications
- o Apple IIgs Technical Note #36, Port Driver Specifications
- o Apple IIgs Technical Note #69, The Ins and Outs of Slot Arbitration
- o Apple IIgs Technical Note #75, BeginUpdate Anomaly

PostScript is a registered trademark of Adobe Systems Incorporated.

Apple IIgs
#94: Packing It In (and Out)

Revised by: Dave Lyons
Written by: C.K. Haun <TR>

May 1992
September 1990

This Technical Note discusses a potential problem with the Miscellaneous Tools routine UnPackBytes.

CHANGES SINCE SEPTEMBER 1990: Noted that the problem detecting the end of the unpack-to buffer near the end of a bank is fixed in System 6.0.

PackBytes and UnPackBytes are handy data compression and expansion routines built into the Apple IIgs System Software. Using them can dramatically reduce the amount of space your application uses on disk or in memory, but you need to understand how these calls work to avoid problems in your applications.

BUFFER SIZE, BUFFER SIZE, BUFF, BUFF, BUFFER SIZE

There are some situations where the Miscellaneous Tools call UnPackBytes does not function as expected and can cause your application to loop infinitely while you're waiting for an unpacking process to finish.

The following packed data and code (in APW assembly) demonstrates the problem. It shows a small routine that unpacks data in two steps, simulating the situation in many applications where an arbitrary amount of data is unpacked in a variable amount of unpacking actions, depending on the results of the last unpack pass.

```
UnPackBuffer    ds    160                ; area to unpack the data to
UnPackBufferPtr dc    i4'UnPackBuffer'    ; pointer to unpacking buffer
UnPackBufferSize ds    2
temp            ds    2

PackedData      dc    h'FFFFFFFF'
EndPackData     anop
PackLength      dc    i2'EndPackData-PackedData' ; how many bytes of packed
data
```

* In packbytes format \$FFFF means '64 repeats of the next byte (\$FF) taken as
* 4 bytes' as described on page 14-39 of Toolbox Reference, so
* this data should unpack into 512 \$FF bytes

* The following code loops infinitely

```
UnPackLoop      lda    #160                ; Unpack buffer size
                 sta    UnPackBufferSize
                 pea    0                    ; return space
                 pushlong #PackedData      ; pointer to packed data
                 pea    2                    ; size of the packed data,
                 ; unpack two bytes
                 ; at a time
                 pushlong #UnPackBufferPtr ; pointer to pointer to
                 ; unpacking buffer
                 pushlong #UnPackBufferSize ; pointer to word with the
```

```

; size of the
; unpacking buffer

_UnPackBytes
pla ; returns 0 bytes unpacked
sta temp
lda PackLength
sec
sbc temp ; subtracting it from our
; known
sta PackLength ; length of packed data
bne UnPackLoop ; this is always be non-zero

```

The problem is in the data and the buffer size. UnPackBytes is being told to unpack two bytes (\$FFFF), which generate 256 bytes of unpacked data, into a 160-byte buffer. Instead of reporting an error with this condition, UnPackBytes instead just does nothing and passes back zero as the returned number of bytes unpacked. If you are relying on the unpacked byte count returned to control your unpacking loop, then you may encounter this problem.

UnPackBytes can be used to unpack in multiple steps, of course, but it cannot unpack a partial record. It cannot unpack 160 bytes of the 256 bytes specified in this record because UnPackBytes does not maintain any state information, so it must unpack full records or do nothing. If the buffer had been 256 bytes, this call would have succeeded.

THE FIX

Fortunately, it's easy to avoid this situation if you know that it can exist. Simply, always supply UnPackBytes with a buffer that is big enough for it to unpack at least two bytes (a flag or count byte and a data byte). The largest value of a flag or count word possible is \$FF, 64 repeats of the next byte taken as four bytes, which generates 256 unpacked bytes. So always give UnPackBytes a 256-byte long output buffer and you should never encounter this problem.

CHECK YOUR CURRENT APPLICATIONS

Please check your current applications to see if you could encounter this problem. One of the most likely places for this error to occur is in applications that process Apple Preferred (file type \$C0, auxiliary type \$0002) pictures. While most pictures currently available are screen-width or less (160 bytes or less per scan line), the Apple Preferred format and QuickDraw II both support pictures that are wider than the current Apple IIgs screen. If someone has created a picture with a PixelsPerScanLine value of 1,280 with a ModeWord of \$0080, it would generate a scan line that was 320 bytes long. If a scan line in this hypothetical picture were all white, for example, the first two bytes of the packed scan line would be \$FFFF, and applications that assume a standard maximum 160 bytes per scan line would not handle this correctly.

BUT THAT'S NOT ALL...

In System Software earlier than 6.0, UnPackBytes has some other buffering problems of which you need to be aware. The size and location of the input buffer (the buffer containing your packed data) can also cause problems.

You can ignore this section if your application requires System 6.0.

NOTE : These problems only occur if you are doing multipass

unpacks. If you always unpack a packed data range in one pass (with one call to UnPackBytes for the whole data set) then you are not affected by these problems, and the restrictions described herein do not apply.

MULTIPASS RESTRICTIONS

When performing a multipass unpack (as described on pp. 14-43..44 of the Apple IIgs Toolbox Reference, Volume 1) the packed data needs to follow two rules.

Rule 1: Your packed data buffer cannot cross a bank boundary.

Rule 2: Your packed data buffer needs to be at least 65 bytes longer than the actual size of the data.

These rules are required by a bug in UnPackBytes. When UnPackBytes begins to unpack a record, it checks the record data to see if there are enough bytes in the current source buffer to unpack the number of bytes requested in the record header (described on pg. 14-39 of the Apple IIgs Toolbox Reference, Volume 1). If there are not enough bytes left for the current record (i.e., the header says to process 63 bytes, and there are only 30 left in the buffer), UnPackBytes returns to the caller. The caller then adjusts the source buffer for the next pass based on the amount of actual bytes unpacked, so the bytes left over from the last pass get processed the next time.

The problem occurs when the partial record is close to the end of a bank. When UnPackBytes checks to see if there is enough data left in the buffer, the check is flawed when the real end of the buffer is near the end of a bank, and a complete copy of the partial record would extend into the next bank. UnPackBytes erroneously thinks that the record is complete, and happily unpacks the remaining actual packed data, plus random information from the next bank. It continues to unpack nonsense data until it fills the unpacking buffer and the number of bytes unpacked returned by the UnPackBytes call is greater than the bufferSize parameter passed as input.

To prevent this bug from occurring, you need to make sure that the buffer for the packed data is at least one record length away from the end of a memory bank. Since the largest packed data record is one flag byte and 64 data bytes, adding 65 bytes to the end of your buffer does the trick. This ensures that your packed data is 65 bytes away from the end.

Following is an example of a safe way to prepare your packed data buffer for multipass unpacking, in APW assembly:

```
* Some data space
myCallBlock dc    i2'2'          ; two parameters
fileRefNum  ds    2              ; file reference number
EOFReturned ds    4              ; file length returned by this call
myIDNumber  ds    2              ; your application memory manager ID
number

* assume that a packed data file is open, and it's a plain packed screen
image, not over 32K
        jsl    $E100A8          ; ask GS/OS for the length of the
data
        dc    i2'$2019'         ; Get_EOF call
        dc    i4'myCallBlock'
```

* Now we need a handle to read it into

```
        pha
        pha                    ; return space
        pea   0                 ; size, high word
        lda   EOFReturned       ; the actual size of the packed data
```

```

sta    actualPackDataSize
clc
adc    #65                ; ask for a handle 65 bytes longer
                        ; than the data

pha
lda    myIDnumber        ; Memory Manager ID for your
                        ; application

pha
pea    $8010             ; attrLocked and attrNoCross
pea    0
pea    0                 ; anywhere
    _NewHandle          ; get the handle

```

Now you have a handle 65 bytes longer than your data that does not cross a bank boundary. You are ready to read in the data and perform a multipass unpack.

PACKBYTES BUFFERS COUNT TOO

PackBytes can also cause you problems if you do not plan for the worst-case situation. Unlike the other toolbox compression routine ACECompress, PackBytes is not guaranteed to shrink the source data. In fact, your data size may actually grow after a PackBytes call.

If you pass a data stream of 64 bytes, all with different values, to PackBytes, PackBytes puts 65 bytes in your output buffer--the 64 original data bytes and the flag byte of \$3F, indicating "64 bytes follow, all different." Unless you preprocess or analyze your data before packing to avoid this situation, make sure your output buffer is large enough to hold the worst case situation, one additional byte generated for every 64 bytes passed to PackBytes for compression.

Further Reference

- o Apple IIgs Toolbox Reference, Volumes 1-3
- o File Type Note for File Type \$C0, Auxiliary Type \$0002, Apple Preferred Format

Apple IIgs
#95: ROM Diagnostic Errors

Written by: Dan Strnad September 1990

This Technical Note describes errors returned by the ROM Diagnostics on Apple IIgs systems.

The Built-In Diagnostics Revealed

The IIgs has a self-test capability in ROM. The self-test is activated by pressing Open-Apple and Option on power up, or Open-Apple, Option, and Reset. During the test, the test number is visible on the bottom of the screen followed by six zeros. After all tests are complete, a continuous 6 KHz one-second beep sounds and the screen displays a System Good message. If any test fails, the screen displays a message of the form System Bad: AABCCDD on the lower left hand side and a staggered AABCCDD on the upper left hand side to help read the error code in the event of a RAM failure. In the event of video failure, the failure code is also sent to the printer port. In the number contained in the error message, AA is the test number that failed and the failure code is embedded in the BB, CC, and DD fields. The complete failure codes for each of the 12 tests are as follows:

Self Test 1: ROM Test

AA = 01
BB = Failed checksum
DD = 01 if the test encountered bad RAM and the error code is a RAM error code similar to the RAM Test error codes

For a failure in ROM, the ROM diagnostics also display RM on the top left hand corner of the screen.

Self Test 2: RAM Test

AA = 02
BB = Bank Number (or \$FF for ADB Tool call error)
CC = Bit(s) failed

Self Test 3: Soft Switches and State Register Test

AA = 03
BB = State Register bit (if any)
CC = Low byte of soft switch address

Self Test 4: RAM Address Test

AA = 04
BB = Failed bank number (or \$FF for ADB Tool call error)
CCDD = Failed address

Self Test 5: Speed Test

AA = 05
BB = 01: Speed stuck slow

02: Speed stuck fast

Self Test 6: Serial Test

AA = 06:
BB = 01: Register R/W
04: Tx Buffer empty status
05: Tx Buffer empty failure
06: All Sent Status fail
07: Rx Char available
08: Bad data

Self Test 7: Clock Test

AA = 07
DD = 01: Fatal error occurred and the test is aborted

Self Test 8: Battery RAM Test

AA = 08
BB = 01: Address test and CC = bad address
02: Non-volatile RAM failed and CC = pattern, DD = address

Self Test 9: Apple Desktop Bus Test

AA = 09
BBCC = Bad checksum
DD = 01: Apple Desktop Bus tools call encountered a fatal error, no checksum computed.

Self Test 10: Shadow Register Test

AA = 0A
BB = 01: Text page 1 fail
02: Text page 2 fail
03: Apple Desktop Bus Tool call error
04: Power On Clear bit error

Self Test 11: Interrupts Test

AA = 0B
BB = 01: VBL interrupt time-out
02: VBL IRQ status fail
03: 1/4 sec interrupt
04: 1/4 sec interrupt
05:
06: VGC IRQ
07: Scan line

Self Test 12: Sound Test

AA = 0C
DD = 01: RAM data error
02: RAM address error
03: Data register failed
04: Control register failed
05: Oscillator interrupt timeout

Further Reference


```
ItemTemplate CloseBut640 = {2,
                             79,265,91,375,
                             buttonItem,
                             CloseStr,
                             0,
                             0,
                             0L};
```

```
ItemTemplate NextBut640 = {3,
                             25,265,37,375,
                             buttonItem,
                             DriveStr,
                             0,
                             0,
                             0L};
```

```
ItemTemplate CancelBut640 = {4,
                               97,265,109,375,
                               buttonItem,
                               CancelStr,
                               0,
                               0,
                               0L};
```

```
ItemTemplate Scroll640 = {5,
                           43,265,55,375,
                           buttonItem,
                           AcceptStr,
                           0,
                           0,
                           0L};
```

```
ItemTemplate Path640 = {6,
                         12,15,24,395,
                         userItem,
                         0L,
                         0,
                         0,
                         0L};
```

```
ItemTemplate Files640 = {7,
                          25,18,107,215,
                          userItem + itemDisable,
                          0L,
                          0,
                          0,
                          0L};
```

```
ItemTemplate Prompt640 = {8,
                           3,15,12,395,
                           statText + itemDisable,
                           0L,
                           0,
                           0,
                           0L};
```

```
/******  
*  
* myDialogHook  
*  
*****/
```

```

pascal void myDialogHook(strip1,strip2)
long strip1;
long strip2;
{
}

/*****
*
* CustomItemDraw
*
*****/

pascal void CustomItemDraw(itemDrawPtr)
Pointer itemDrawPtr;
{
static unsigned int flag, dbr;          /* result, data bank register value */
byte      StringCount;
char      *ItemPascalString;
Word      ItemFileType;
Long      ItemAuxType;
Rect      *TheItemRectPtr;
MemRec    *TheMemRecPtr;
CtlRecHndl TheSFListControlHndl;
Point     MyOldPenPos,
          MyNewPenPos;

static char FileString[] = "xxxx YYYYYYYY ";

/* save our data bank and set current to global page */
dbr = SaveDB();
/* Get the Rect from High on the Stack */
TheItemRectPtr = (Rect *)(((long *)((long)&itemDrawPtr)+ 36L));

GetPen(&MyOldPenPos);                    /* save old pen position */
MyNewPenPos.h = TheItemRectPtr->h1 + 5;  /* Set our pen position */
MyNewPenPos.v = TheItemRectPtr->v2 -2;
MoveTo(MyNewPenPos);                    /* relocate the pen */

/* get our member record; this is just to reveal where it is on the stack */
TheMemRecPtr = (MemRec *)(((long *)((long)&itemDrawPtr)+ 32L));

/* get the list cntrol handle; ditto */
TheSFListControlHndl = (CtlRecHndl)(((long *)((long)&itemDrawPtr)+ 28L));

StringCount = (byte) *itemDrawPtr;      /* get the string length */
ItemPascalString = itemDrawPtr;        /* set our user string */
ItemFileType = *(Word *) (itemDrawPtr+StringCount+1L); /* get our FileType */
ItemAuxType = *(Long *) (itemDrawPtr+StringCount+3L); /* get our AuxType */

/* format for display */
sprintf(FileString, "%.4x-%.8lx ",ItemFileType,ItemAuxType);
c2pstr(FileString);                    /* turn it into a P string */
DrawString(FileString);                /* Draw it */
DrawString(ItemPascalString);          /* catenate File name to the other info */
FrameRect(TheItemRectPtr);
MoveTo(MyOldPenPos);                  /* return the pen to starting position */
RestoreDB(dbr);                       /* restore our data bank */

}

/*****

```

```

*
* ChooseFolder
*
* presents user with dialog to select folder to show/set privileges of
*
*****/

void    SomeProc()
{
DialogTemplate GetDialog640;

GetDialog640.dtBoundsRect.v1 = 0;
GetDialog640.dtBoundsRect.h1 = 0;
GetDialog640.dtBoundsRect.v2 = 114;
GetDialog640.dtBoundsRect.h2 = 400;
GetDialog640.dtVisible = -1;
GetDialog640.dtRefCon = 0L;
GetDialog640.dtItemList[0] = &OpenBut640;
GetDialog640.dtItemList[1] = &CloseBut640;
GetDialog640.dtItemList[2] = &NextBut640;
GetDialog640.dtItemList[3] = &CancelBut640;
GetDialog640.dtItemList[4] = &Scroll640;
GetDialog640.dtItemList[5] = &Path640;
GetDialog640.dtItemList[6] = &Files640;
GetDialog640.dtItemList[7] = &Prompt640;
GetDialog640.dtItemList[8] = 0L;

SFPGetFile2( /* user selection of folder to get/set privs of */
             120, 53,
             CustomItemDraw,
             refIsPointer,
             prompt,
             0L,
             0L,
             &GetDialog640,
             myDialogHook,
             &myReply
);

```

Further Reference

-
- o Apple IIgs Toolbox Reference, Volumes 1 & 3
 - o DTS Apple II Sample Code #18, AccessPriv

Apple IIgs
#97: Picture Comments and Printing

Written by: Matt Deatherage, Suki Lee & Ben Koning

November 1990

This Technical Note discusses QuickDraw Auxiliary picture comments and how they can be used to help control the printing process.

What's a Picture Comment?

Picture comments are a way in which extra information beyond normal QuickDraw II calls can be embedded in a QuickDraw II picture. Comments can contain virtually anything; they consist of a length, a handle containing the comment and a "kind" that identifies the general type of information in the comment. Picture comment kinds less than or equal to 256 (\$100) are reserved for Apple Computer, Inc.

For comments to have any significance, there must be a way that a routine can take special action on them. One of the standard bottleneck procedures is called every time a picture comment is encountered, and it is passed the picture comment's kind, size, and handle on QuickDraw II's direct page. You can insert the address of a custom picture comment handler into the grafProcs field of a grafPort as described in Apple IIgs Technical Note #34, Low-Level QuickDraw II Routines. If no custom comment handler is present in the grafPort, the system calls its own StdComment routine, which ignores all comments.

The current picture comment handling routine (either a custom one or the system's default one) is called whenever a picture comment is generated (with the QuickDraw Auxiliary call PicComment) or played back from a picture (from within DrawPicture). Since the picture comment handling procedure is called when the comment is created, a picture does not have to be open for this facility to work.

Picture comments are ideal ways for applications to pass information to printer drivers as they are generated through toolbox calls and are easily accessible to any desktop program. If the printer driver is printing in immediate mode, it can intercept and act on the picture comment when it is generated. If the printer driver is printing in deferred mode and recording page images with QuickDraw II pictures, it can intercept and act on the picture comment when the picture is played back.

Apple's ImageWriter, ImageWriter LQ and LaserWriter drivers (from System Software 5.0.3) all support various kinds of picture comments for controlling printed output. Applications are encouraged to use these picture comments for finer control over printing. Authors of printer drivers are encouraged to act on these picture comments where appropriate, so applications which use them achieve similar results across printing platforms.

The LaserWriter Driver's Picture Comments

Version 2.2 and later of the LaserWriter driver support the following five PostScript picture comments:

Name	Kind	Size	Handle

PostScriptBegin	190	0	NIL
PostScriptEnd	191	0	NIL
PostScriptHandle	192	-	PostScript data
PostScriptFile	193	-	PostScript path name
TextIsPostScript	194	0	NIL

Table 1-PostScript Picture Comments

The print loop must be completed normally with or without any PostScript picture comments that are included. PostScript transmission must begin with the PostScriptBegin picture comment and end with the PostScriptEnd picture comment. Never nest PostScriptBegin and PostScriptEnd picture comments.

The PostScriptHandle picture comment takes a handle containing PostScript commands (in the form of ASCII data) and sends it to the LaserWriter. The size field must contain the size of the handle.

The PostScriptFile picture comment takes a handle containing the pathname of a disk file containing PostScript commands. The size field must contain the size of the pathname.

The TextIsPostScript picture comment takes text drawn through the QuickDraw II StdText bottleneck and sends it to the LaserWriter as PostScript. This picture comment has the effect, from the application's point of view, of interpreting all strings passed to DrawString and similar calls as PostScript. This picture comment is specific to LaserWriters (idev = \$0003). Other drivers do not implement this picture comment; therefore, text drawn through QuickDraw II is simply printed-it is neither interpreted as PostScript nor ignored.

The driver does not check for PostScript errors, so the data sent to the LaserWriter must be correct. Always terminate PostScript text with a carriage return character. The transformation the driver uses flips text and prints it upside down on the page. Applications should set their own transformation matrices to serve their needs. Never use the LaserWriter's userdict-define a local dictionary for your application's use. Never use exitserver, initgraphics, grestoreall, erasepage, or showpage PostScript commands, as these commands can alter the driver's environment.

See Chapter 3 of the LaserWriter Reference Manual for some examples of how to use picture comments.

The ImageWriter Driver's Picture Comments

ImageWriter driver version 4.0 and later uses three picture comments to control alternate color selection:

Name	Kind	Size	Handle
Reserved	250	-	Reserved
FillColorTable	251	42	See below
ChangeSCBs	252	14	See below

The structure passed in the handle to FillColorTable looks like the following:

version	word	must be zero
signature	word	must be \$A55A
tableno	word	the color table to be modified (0-15)
table	32 bytes	The new color values for the color table
reserved	long	must be zero

The structure passed in the handle to ChangeSCBs looks like the following:

version	word	must be zero
signature	word	must be \$A55A
Y1	word	line number of first SCB to change (from zero to rPage.Y2)
Y2	word	line number of last SCB to change (from zero to rPage.Y2)
SCBvalue	word	the new scan line control byte
reserved	long	must be zero

PrOpenPage reinitializes the printing grafPort, so these picture comments should be used immediately after PrOpenPage to set custom colors. Having lines with both 320- and 640-mode is not recommended and will probably not be supported in the future.

The ImageWriter driver uses picture comment 250 to internally mark the end of a page. Applications must not use this picture comment; use PrClosePage to end a page's definition.

Further Reference

-
- o Apple IIgs Toolbox Reference, Volumes 1-3
 - o Apple IIgs Technical Note #34, Low-Level QuickDraw II Routines
 - o Apple IIgs Technical Note #35, Printer Driver Specifications
 - o Apple IIgs Technical Note #93, Compatible Printing
 - o d e v e l o p, October 1990, Issue 4, "Driving to Print"

PostScript is a registered trademark of Adobe Systems, Incorporated.

Apple IIgs

#98: Aren't Windows A Pane?

Revised by: Dave Lyons

May 1992

Written by: Dave Lyons

January 1991

This Technical Note describes interesting Window Manager things.

CHANGES SINCE JANUARY 1991: Noted that in System 6.0 it's safe to use Window color table resources. Added a section on changing the desktop pattern or picture.

CHANGING THE DESKTOP PATTERN OR PICTURE

The best way to set a new desktop pattern or picture is as follows. This works with the Finder and other desktop applications.

1. Use MessageCenter to delete message 2, the desktop message. (If there wasn't one, that's fine--there still isn't.)
2. Use MessageCenter to create a new message 2, containing the pattern or picture you want (see the Window Manager chapter of Apple IIgs Toolbox Reference, Volume 3).
3. Call Desktop (in the Window Manager) with a deskTopOp of 8 and a dtParam of \$00000000. This notifies any part of the system that cares (such as the Finder) that there is a new desktop pattern.
4. Call Desktop with a deskTopOp of 4 and a dtParam of \$00000000 and keep the result.
5. Call Desktop with a deskTopOp of 5 and use the result from step 4 as dtParam. This sets the desktop pattern to what it already is, forcing the desktop to redraw (this works whether a pattern, picture, or pointer to desktop-drawing routine is involved).

A WARNING ABOUT WINDOW COLOR TABLE HANDLES AND RESOURCES

The System 6.0 Window Manager fixes the problem described below. If your application requires System 6, you can safely ignore this section.

All versions of the Window Manager that support window color tables specified as handles or resources, up to and including System Software 5.0.4, work unreliably when a standard window's color table is supplied by handle or resource ID.

The problem is not immediately obvious; only one bit of memory is accidentally cleared, but the address is unpredictable in advance. (When unlocking the color table handle, the standard window definition procedure attempts to unlock the handle manually by turning off bit 15 of word offset +4 in the master pointer record. But it gets the high and low words of the handle reversed and usually turns off bit 15 of the word at offset \$80E4 or \$80E5 in some bank of RAM determined by the low byte of the handle.)

The solution is to avoid supplying color table handles or resource IDs to the Window Manager. Supply color table pointers instead. You can get a color table pointer from a color table resource ID by calling LoadResource on the color table resource, locking the handle and dereferencing it. Memory is less fragmented if color table resources used in this way are marked as attrFixed.

One method is to put the window color table pointer into the window template before calling `NewWindow2`. If you are creating the window from an `rWindParam1` resource, you need to use `LoadResource` to get the template into RAM so that you can stuff the color table pointer into the template. (Be sure to change the `moreFlags` field to indicate that the color table is a pointer, if the template indicates it's a resource.) After you create the window with `NewWindow2` (by handle), use `ReleaseResource` to release the `rWindParam1` resource.

Another method is to create the window as invisible and pass the window color table pointer to `SetFrameColor` before calling `ShowWindow`.

Further Reference

- o Apple IIgs Toolbox Reference, Volumes 2-3

Apple IIgs
#99: Supplemental Scrap Types

Revised by: Dave Lyons
Written by: Matt Deatherage & Matthew Reimer

May 1992
January 1991

This Technical Note describes public scrap types.

CHANGES SINCE MARCH 1991: Added information on Scrap Type \$8003 (Resource Reference Scrap); added a cross-reference to HyperCard IIgs Technical Note #3.

The Apple IIgs Toolbox Reference lists only two known scrap types--text (\$0000) and pictures (\$0001). Other assigned scrap types are documented in this Note. The format used to describe the scraps is similar to that used in File Type Notes, where the offsets, given in the form (+xxx), determine the offset from the beginning of the scrap handle.

SAMPLED SOUND SCRAP (TYPE: \$0002)

The following describes the Sampled Sound scrap format. It consists of a ten-byte header followed by the sample data bytes. This format is identical to the sampled sound resource format.

Format	(+000)	Word	This must always be zero.
Wave Size	(+002)	Word	Sample size in pages (256 bytes per page). For example, an 8K sample takes 32 pages; a 128K sample requires \$200 pages.
Rel Pitch	(+004)	Word	The high byte of this word is a semitone value; the low byte is a fractional semitone. These values are used to tune the sample to correct pitch. (See HyperCard IIgs Technical Note #3, Tuning Sampled Sounds.)
Stereo	(+006)	Word	The output channel for this sound is in the low nibble of this word.
Sample rate	(+008)	Word	The sampling rate of the sound, in Hertz (Hz).
Sound	(+010)	Bytes	The sampled sound data. The bytes are all 8-bit samples. The sample starts here and continues until the end of the scrap.

TEXTEDIT STYLE SCRAP (TYPE: \$0064)

The TextEdit Style Scrap format is the same as the TEFormat structure defined in Volume 3 of the Apple IIgs Toolbox Reference, which is also the same as the rStyleBlock resource format defined in the same volume.

ICON SCRAP (TYPE: \$4945)

The Icon scrap format is the same as the format for Finder Icon Data records, documented in detail in the File Type Note for File Type \$CA, Finder Icon Files. If there is more than one Icon Data record in a scrap, they are concatenated together with no intervening space.

MASK SCRAP (TYPE: \$8001)

The Mask scrap format is exactly the same as the PICT scrap (\$0001) format, except that the pixel image the picture draws contains only zeroes and ones. When drawn, this picture creates a mask. The mask has zeroes where the image can be seen through the mask, and ones where the mask does not allow the picture through. When pasting a Mask scrap, initialize the destination bitmap to zero and draw the picture.

You can create the mask image by using regular QuickDraw II calls (using ovals, rectangles, etc.) or you can create it independently and include it with PaintPixels or other pixel map manipulation routines.

COLOR TABLE SCRAP (TYPE: \$8002)

The following describes the Color Table scrap format. The scrap contains color tables so that applications can keep custom colors with pictures copied to the clipboard. The scrap has the same format as the Apple Preferred Format picture PALETTES block:

NumColorTables	(+000)	Word	The count of the number of color tables in the scrap
ColorTableArray	(+002)	32 Bytes	The color tables for the scrap. There are NumColorTables of them, each 32 bytes long.

RESOURCE REFERENCE SCRAP (TYPE: \$8003)

The Resource Reference scrap is designed to allow resource editors to exchange resource data through an external scrap file using the Scrap Manager.

resScrapType	(+000)	Word	Type of resource (within the resScrapPath file)
resScrapID	(+002)	Long	ID of resource (within the resScrapPath file)
resScrapPath	(+006)	WString	Full GS/OS class-one pathname to an extended file containing the specified resource.

If the specified resource contains references to other resources (for example, an rWindParam1 resource with a title string, control list, control templates, etc.), all the referenced resources must be present in the resScrapPath file.

It is the responsibility of the application using this scrap to handle resource ID conflicts that might arise from a Paste operation. The application should not modify or destroy the resScrapPath file.

Further Reference

- o Apple IIgs Toolbox Reference
- o HyperCard IIgs Technical Note #3, Tuning Sampled Sounds
- o File Type Note for file type \$CA, all auxiliary types, Finder Icons File
- o File Type Note for file type \$C0, auxiliary type \$0002, Apple Preferred Format

maybe a later one. This mechanism is your primary defense against old system software--by requiring the latest tool versions in your application, you are notified by the Tool Locator early in your program if the system has the latest system software installed or not.

Note that ROM 1 and ROM 3 have different version numbers for seven tools under 5.0.4--QuickDraw II, the Scheduler, ADB, SANE, Integer Math, Text Tools and the List Manager. In each case, the ROM 01 version is lower and should be used in your LoadOneTool, LoadTools or StartUpTools calls. The current revision of Apple IIgs System Software is 6.0. Assuming a correct installation, requiring QuickDraw 3.7 in effect requires System Software 6.0, although you may check the system's rVersion resource in the system resource file if you require more detailed information about the system software version.

System Tool Set Versions

Number	Tool	ROM 1	ROM 3
1	Tool Locator	\$0301	\$0301
2	Memory Manager	\$0302	\$0302
3	Misc Tools	\$0302	\$0302
4	QuickDraw II	\$0307	\$0307
5	Desk Manager	\$0304	\$0304
6	Event Manager	\$0301	\$0301
7	Scheduler	\$0300	\$0300
8	Sound Tools	\$0303	\$0303
9	ADB	\$0300	\$0300
10	SANE	\$0300	\$0300
11	Integer Math	\$0300	\$0300
12	Text Tools	\$0300	\$0300
13	[used internally]	\$0300	\$0300
14	Window Manager	\$0303	\$0303
15	Menu Manager	\$0303	\$0303
16	Control Manager	\$0303	\$0303
17	[System Loader]	\$0400	\$0400
18	QuickDraw II Aux	\$0304	\$0304
19	Print Manager	\$0301	\$0301
20	Line Edit	\$0303	\$0303
21	Dialog Manager	\$0304	\$0304
22	Scrap Manager	\$0301	\$0301
23	Standard File	\$0303	\$0303
25	Note Synthesizer	\$0104	\$0104
26	Note Sequencer	\$0104	\$0104
27	Font Manager	\$0303	\$0303
28	List Manager	\$0303	\$0303
29	ACE	\$0103	\$0103
30	Resource Manager	\$0102	\$0102
32	MIDI Tools	\$0103	\$0103
33	Video Overlay	\$0103	\$0103
34	Text Edit	\$0103	\$0103
35	MIDI Synth	\$0100	\$0100
38	Media Control	\$0100	\$0100

Toolbox Driver Version Numbers

Driver	Version
ImageWriter II	4.2
ImageWriter LQ	4.2
LaserWriter	3.2
StyleWriter	1.0
Epson	2.0

Printer Port Driver	2.1
Modem Port Driver	2.1
Parallel Card Port Driver	2.0
AppleTalk Port Driver	3.0
Pioneer 4200 (MC)	1.0
Pioneer 2000 (MC)	1.0
Apple CD SC (MC)	1.0

GS/OS Version Numbers

Component	Version
GS/OS	4.1
ProDOS FST	4.1
AppleShare FST	4.0
High Sierra FST	4.0
Character FST	4.0
DOS 3.3 FST	1.2
HFS FST	1.0
Pascal FST	1.0
AFP Driver	4.0
Apple II RAMCard driver	1.0
AppleDisk 3.5 Driver	5.3
AppleDisk 5.25 Driver	2.5
AppleTalk Main Driver	4.0
Console Driver	3.2
RPM Driver	4.0
SCSI CD Driver	6.0
SCSI HD Driver	6.0
SCSI Scanner Driver	6.0
SCSI Tape Driver	6.0
UniDisk 3.5 Driver	3.0

Control Panel Version Numbers

CDev	Version
AppleShare	2.0
Direct Connect Printer	1.1
FolderPriv	1.0
General	2.0
Keyboard	1.1
Media Control	1.1
MIDI	1.0
Modem Port	1.1
Monitor	1.1
Network Printer Namer	1.0
Network Printer Chooser	1.0
Network	1.0
Printer Port	1.1
RAM	1.1
SetStart	1.0
Slots	1.2
Sound	2.0
Time	2.0

Further Reference

-
- o Apple IIgs Toolbox Reference
 - o GS/OS Reference

- o GS/OS Technical Note #1, Contents of System Disk and System Tools
- o File Type Note for File Type \$C7, Control Panel Devices

Apple IIgs
#101: Patching the Toolbox

Revised by: Dave Lyons
Written by: Dave Lyons

May 1992
May 1991

This Technical Note presents guidelines on when and how to patch Apple IIgs Toolbox functions.

CHANGES SINCE MAY 1991: Added a note about patching the Tool Locator and Desk Manager, and corrected a spelling error.

INTRODUCTION

There is normally no need to patch the toolbox; avoid patching whenever you can. If you must patch a toolbox function, be sure to have a good understanding of the call you're patching and how it interacts with the whole system.

No toolbox patch is risk-free. Future versions of the toolbox could change in ways that make your patch less useful. (For example, if you patched NewControl to have some global effect on controls being created, your patch became less useful when NewControl2 was introduced in System Software 5.0.)

For better compatibility, patch with care! If any parameters passed are outside the range that was allowed when you wrote your patch, just pass the call straight through; the new toolbox probably knows something your patch doesn't.

PATCHING THE TOOLBOX FROM AN APPLICATION

An application can easily patch a function for the duration of that application.

After starting up the tools, construct a Function Pointer Table (FPT) the same size as the existing FPT (call GetTSPtr and examine the first word of the table; multiply it by four to get the size of the FPT in bytes). The first longword of your FPT is the number of functions in the tool set; do not hard-code this value! Get it from the existing FPT on the fly. Fill the rest of your FPT with zeroes, except for the functions you want to patch. You must always patch the BootInit function (the first function) to return no error. Remember that the function pointer values are one less than the addresses of your replacement functions.

On exit, when you call TLShutDown your patch will be automatically removed. (If you're using ShutDownTools, you should call MMShutDown and TLShutDown after you call ShutDownTools.)

Note : In the description of SetTSPtr on page 24-19 of Apple IIgs Toolbox Reference, Volume 2, there are several references to the TPT. Keep in mind that the TPT is the Toolset Pointer Table, not the Function Table Pointer you pass to SetTSPtr. While SetTSPtr copies the TPT to RAM if necessary, it does not make a copy of the FPT. After you call SetTSPtr, the FPT you passed is being

used, and any zero values in your table were filled in.

PATCHING THE TOOLBOX FROM A DESK ACCESSORY OR SETUP FILE

A permanent initialization file or Desk Accessory can patch toolbox functions at boot time by constructing an FPT for SetTSPtr, as described for an application, but there is an extra step to make the patch "stick."

Call LoadOneTool and then SetTSPtr; then call SetDefaultTPT (see Apple IIgs Toolbox Reference Volume 3, page 51-16).

It is not safe to call SetDefaultTPT while an application is running (temporary application patches would be made permanent, and later the application would go away). Since there are desk accessories that install other desk accessories while applications are running, desk accessory that wants to install a tool patch should make the class-one GS/OS GetName call; if the null string is returned, no application is executing yet, so it is safe to make the patch. (Otherwise the desk accessory should ask the user to put the desk accessory file in the System:Desk.Accs folder and restart the system.)

PATCHING THE TOOL LOCATOR OR DESK MANAGER

On ROM 3 systems, the SetTSPtr call treats toolsets 1 (Tool Locator) and 5 (Desk Manager) specially, for compatibility with system software versions earlier than 5.0.

You must pass a systemOrUser value of \$0001 (not \$0000) when patching one of these toolsets, or the SetTSPtr call will have no effect. Passing this special systemOrUser value works for other ROM versions, too--you don't have to check the ROM version.

AVOID TAIL PATCHING

The best kind of patch is a pre-patch or head patch: it does some extra work and then jumps to the original function (as found in the FPT before applying the patch). Make sure the A, X, and Y registers contain the same values when you jump to the original function as they did when the patch got control.

A "tail patch" which calls the original function and then regains control is much more of a compatibility risk, because there are several instances where System Software patches examine return addresses to fix problems in large toolbox calls which call small ones (by patching the small one to realize it's being called from the big one, many K of RAM remain available to your application).

If you tail patch a function which the system already patched, you may prevent the toolbox from working correctly.

PATCHING THE TOOL DISPATCHER

If you need to patch a large number of functions, especially for a general purpose utility like a debugger, it may make more sense to patch the tool dispatcher vectors instead of patching individual functions. See Apple IIgs Technical Note #87, Patching the Tool Dispatcher.

Further Reference

- o Apple IIgs Toolbox Reference
- o Apple IIgs Technical Note #87, Patching the Tool Dispatcher

Apple IIgs
#102: Various Vectors

Revised by: Dave Lyons
Written by: Dave Lyons

May 1992
December 1991

This Technical Note describes system vectors that are not fully described in other documentation.

CHANGES SINCE DECEMBER 1991: Added information about the TOBRAMSETUP vector.

THE TOBRAMSETUP VECTOR

The TOBRAMSETUP vector is documented in Appendix D of the Apple IIgs Firmware Reference. Two clarifications are needed:

- o TOBRAMSETUP must be called in 8-bit native mode (SEP #30).
- o Before System 6.0, TOBRAMSETUP required that the Bank register be \$00 (bad things would happen if it was not). This requirement is gone in 6.0.

THE MOVE_INFO VECTOR

MOVE_INFO is a flexible, low-overhead data transfer routine. It can transfer buffer-to-buffer, buffer-to-location, location-to-buffer, and buffer-to-buffer reversing the order of the bytes.

Apple IIgs GS/OS Device Driver Reference tells you how to call MOVE_INFO from a GS/OS driver environment (JSL to \$01FC70), but this requires the language-card RAM to be banked in correctly.

Another vector points to the same routine: \$E10200. If you aren't a GS/OS device driver, it is more convenient to JSL to \$E10200, because you don't have to worry about banking in the \$01FCxx vectors. The \$E10200 vector is available whenever GS/OS is active, under System Software 5.0 or later.

THE DYN_SLOT_ARBITER AND SET_SYS_SPEED VECTORS

Two other GS/OS System Service vectors are duplicated in bank \$E1: SET_SYS_SPEED (\$E10204) and DYN_SLOT_ARBITER (\$E10208). Like MOVE_INFO, these are available when GS/OS is active under System Software 5.0 or later.

Further Reference

- o Apple IIgs GS/OS Device Driver Reference
- o Apple IIgs Firmware Reference

Apple IIgs
#103: Inline Procedure Name Format

Modified by: Matt Deatherage
Written by: Dave Lyons

May 1992
December 1991

This Technical Note describes a simple format for imbedding procedure names in object code, for use by debugging utilities.

CHANGES SINCE DECEMBER 1991: Changed &syscnt to &SYSCNT so it works with the CASE ON APW directive. Clarified the possible addition of parameters after the Pascal string.

GSBug 1.5b18 and later support a simple convention for including procedure names inline in the object code, for debugging purposes.

INLINE NAME FORMAT

```
82 xx xx          brl  pastName
71 77            dc.w  $7771
nn xx xx xx xx... str  'the name string'
                    pastName ...
```

That is, an imbedded name is a BRL around a signature word and a Pascal string. The name string can theoretically be up to 255 characters long, but in practice only short names are useful. For example, GSBug displays only the first 15 characters of a name when it is encountered, and only the first 11 when it appears as the operand of a JSR or JSL instruction.

Names in this format always start with a BRL, not a BRA or JMP. Signature word values other than \$7771 are reserved for future definition, and more information may be added after the Pascal string.

Be careful what you name!

Be careful not to name something important--like a table, or a label from which you compute other addresses. The extra bytes generated by the inline name would mess up your calculations. If you name a heartbeat task, out-of-memory queue routine, or other construction that needs a special header, be sure to put the name where the executable code starts, not at the beginning of the header.

APW ASSEMBLY MACRO

The following macro is for the APW assembler. If you equate DebugSymbols to zero, the macro generates no object code. If DebugSymbols is nonzero, the macro generates an inline name corresponding to its label. Use the name macro anywhere you would use a label. For example:

```
DebugSymbols    GEQU 1
...
CountItems     name
```

The macro:

```

MACRO
&lab name
&lab anop
    aif DebugSymbols=0, .pastName
    brl pastName&SYSCNT
    dc i'$7771'
    dc il'L:&lab',c'&lab'
pastName&SYSCNT anop
.pastName
MEND

```

MPW IIgs Assembly Macros

The following macros are for the MPW IIgs assembler. If you equate DebugSymbols to zero, the macros generate no object code. If DebugSymbols is nonzero, the macros generate inline names corresponding to their labels.

Use the name macro anywhere you would use a label. Use the procname macro in place of a proc directive, at the beginning of a procedure. For example:

```

DebugSymbols    equ 1
...
CountItems     name
TaskLoop       procname

```

The macros:

```

&lab           macro
&lab           name
&lab
                if DebugSymbols<>0 then
                brl @pastName
                lclc &olds
&olds          setc &setting('string')
                string asis
                dc.w $7771
                dc.b &len(&lab), '&lab'
                string &olds
@pastName
                endif
                mend

```

* You can use procname instead of proc

```

&lab           macro
&lab           procname &x
&lab           proc      &x
                if DebugSymbols<>0 then
                brl @pastName
                lclc &olds
&olds          setc &setting('string')
                string asis
                dc.w $7771
                dc.b &len(&lab), '&lab'
                string &olds
@pastName
                endif
                mend

```

If you write a utility that recognizes inline procedure names in this format, check for a signature word of \$777x, not specifically \$7771. This allows more information to be added to the format later (a signature of \$7772 could mean there is a Pascal string followed by parameter-passing information, for example).

Apple IIgs
#104: Font Manager Fundamentals

This Technical Note discusses information and philosophy of that typographical toolset, the Font Manager.

FixFontMenu only works once per FMStartUp

You may have noticed that none of the Font Manager calls that translate font family numbers to menu item IDs (or vice-versa) require a menu ID as a parameter. That's because the Font Manager was designed with the idea that an application would only need one font menu, so it keeps one correspondence in private static storage.

This means that once someone has called FixFontMenu, any later FixFontMenu call during that Font Manager session will destroy the results of the first one, unless all the parameters are identical. The Font Manager doesn't remove the font menu items, but it does not return the correct results from FamNum2ItemID or ItemID2FamNum.

This means if you're a new desk accessory, the Font Manager can't help you create a font menu--attempting to use FixFontMenu will make any application font menu useless. You can use Font Manager routines such as CountFamilies, FindFamily and GetFamInfo to obtain all the information necessary to build your own font menu (or font choosing dialog box, for that matter--but if you create a dialog for an NDA, remember that it has to fit in 320 mode also).

Font styling requires QuickDraw Auxiliary

The Font Manager can't create fonts with outline, shadow or italic styles unless QuickDraw Auxiliary (tool set #18) is present and started. These facts are mentioned in pieces other places, but not in one place--if you want normal Font Manager operations, you must load and start QuickDraw Auxiliary.

Further Reference

-
- o Apple IIgs Toolbox Reference, Volumes 1-3

Apple IIgs
#105: We Interrupt This CPU...

Written by: Matt Deatherage

May 1992

This Technical Note supplements the discussion of how interrupts generally work (or don't work) on the Apple IIgs found in the Apple IIgs Firmware Reference. It also discusses how to patch into the interrupt chain and when not to use software interrupts.

THIS NOTE IS A SUPPLEMENT

That's right, a supplement. This is not the definitive, end-all discussion of interrupts on the Apple IIgs. Most of the information you need to know is available, and has been for several years, in the Apple IIgs Firmware Reference. If you're going to write an interrupt routine, you need to read Chapter 6 of the Firmware Reference.

No excuses. If you don't have the book, buy it or borrow it. People who use your software don't want to hear a sad story about how you wanted to spend the money on a couple of CDs instead of preventing their machine from crashing.

If you haven't read Chapter 6 of the Firmware Reference, do so before continuing; the rest of this Note will make much more sense if you're familiar with the material covered in that chapter.

A NOTE ABOUT TIMING

There are lots of times listed in this Note, concerning how fast certain kinds of interrupts must be serviced before they're lost. Please remember that all times listed are ideal times--actual times are likely to be shorter. For example, a maximum response time of a millisecond means you have one millisecond from the time the peripheral asserts the /IRQ line until the interrupt must be serviced. If interrupts are disabled for the first 750 microseconds (us) of that, then your maximum response time is 250 us. This is why we constantly remind programmers to keep interrupts disabled for absolutely the shortest time possible. Also, all times reflecting serial or AppleTalk interrupts already take into account the serial chip's internal 3-byte buffer.

SO WHAT THE HECK ARE ALL THOSE VECTORS?

At first, looking at all those various vectors seems pretty darned intimidating. However, the structure becomes clearer when you think about interrupt priority.

Some microprocessors allow interrupt requests to have priorities--higher priority interrupts can interrupt lower priority ones. The 65816 doesn't have this capability, so the best the Apple IIgs can do is check possible interrupt sources in highest-priority-first order. For example, AppleTalk interrupts must always be processed extremely quickly--from the time an AppleTalk interrupt is asserted, someone must read the data from the SCC within a maximum of 104.167 us or data can be lost. That's not very much time at all, especially considering that the system may have interrupts disabled, or may be

running at 1 MHz speed when the interrupt fires.

Serial interrupts are next--at 19,200 baud, there's a maximum of 1.094 milliseconds to read data before it's lost. (Multiplication shows that 38,400 baud has a maximum of 547 us, and 57,600 baud has a maximum delay of 273.5 us. Not much at all.)

You'd hope the Interrupt Manager in ROM would be smart enough to service AppleTalk interrupts first and serial interrupts next, and in fact that's what it does. In fact, it services them so fast that not all the system information is saved before checking the hardware and dispatching (if necessary) to the IRQ.APTALK or IRQ.SERIAL vectors. See Apple IIgs Technical Note #24 for more information on which system state information isn't saved before calling those vectors.

The list of interrupt priorities is on page 180 of the Firmware Reference. What's not clear from any description of interrupt handling is that each internal interrupt source's vector is only called if the Interrupt Manager determines it is the source of the interrupt. For example, the IRQ.DSKACC vector is not called unless the user pressed Command-Control-Esc to generate the interrupt. This insures that external interrupt handlers for slot-based peripherals are dispatched to as quickly as possible--if each vectored routine had to determine interrupt ownership, every interrupt would have significantly more overhead.

There are two additions to the priority list in the Firmware Reference--the first is also an exception to the "interrupt handlers don't have to identify the interrupt" rule. On ROM 3 machines only, vector \$E1021C (IRQ.MIDI) gets control immediately after determining the interrupt isn't an AppleTalk interrupt. MIDI data can come in so quickly that it needs higher priority than serial interrupts. However, to improve performance, routines called through this vector must return as fast as possible (faster would be better) to avoid delaying interrupts further down the chain, like serial interrupts. Also note that this vector doesn't exist on ROM 1.

The second addition is to the final priority, simply defined as "external slot." The documentation doesn't clearly indicate how this works--it kind of implies this is just calling IRQ.OTHER. In fact, if no IRQ.OTHER routine claims the interrupt, the system does some voodoo magic to switch to emulation mode and jumps through the vector at \$03FE, just like all previous Apple II models. And just like in older systems, whatever code is pointed to by \$03FE must end with an RTI instruction. This behavior is preserved for compatibility, although it is the slowest interrupt response available on the IIgs.

GETTING CONTROL IN TIME

Passing control to external handlers isn't always quick enough for some people. If you're writing a telecommunications program, for example, you have no more than 1.094 ms from the time a character is received to get it out of the SCC or you'll lose data at 19,200 baud.

The Interrupt Manager is a very tight piece of code--if it were running in RAM and the system was temporarily slowed down to 1 MHz, there would only be room for about two more instructions before AppleTalk would lose data. Since AppleTalk has to be serviced within 104.2 us (as discussed previously), and since IRQ.SERIAL is called as quickly as possible after IRQ.APTALK (the only delay is if you're on ROM 3 and a non-trivial MIDI interrupt handler is installed), patching in at IRQ.SERIAL poses no problems for most high-speed communications, even up to 57,600 baud. In other words, it's not necessary to patch any vector other than IRQ.SERIAL to achieve the results you want.

The problem comes when you have external communications hardware--making it through the internal interrupt chain is too slow if your external communications hardware has the same kinds of limitation the SCC does (namely, a 3-byte internal buffer). External vectors are only called after all the internal sources verify it's not their interrupt, and by that time your card may have lost data.

PATCHING THE MAIN INTERRUPT VECTOR

In these cases, where there is no possible way to service an interrupt in time through the Interrupt Manager's normal priority chain, and in these cases only, it's acceptable to patch out the main interrupt vector at \$E10010 (preferably using GetVector and SetVector with reference number \$0004). But even then, there are rules to follow.

1. You should duplicate the functionality of the main interrupt vector exactly until the point where you must gain control or lose data. For example, if your card requires that you service interrupts within a millisecond or lose data, AppleTalk interrupts still have higher priority over your interrupts because AppleTalk interrupts must be serviced within 104 us. In this example case, your code should duplicate the functionality of the Interrupt Manager up through and including the call to IRQ.APTALK, and then (and only then) call your interrupt handler, where you handle the interrupt if it's yours and pass control to the rest of the interrupt chain if it's not.
2. You should only service your interrupts before AppleTalk if your interrupts require servicing in less than 104 us. If they don't, give AppleTalk first shot. If they do, you must clearly inform the user, both in documentation and on the screen, that if they proceed with this function network services may be interrupted, and that they may have to restart the system to restore them. Users must also have the option to back out and cancel at this point. No, this isn't a pleasant message to deliver, but it's much nicer than to completely disconnect AppleTalk and lock up the system if it was booted from a server.
3. You should only patch out the main interrupt vector when absolutely necessary. For example, if you're communicating with hardware that runs at multiple speeds and only the highest speed generates interrupts that require patching the main vector, you should not be patching the main vector when not using that highest speed. For telecommunication programs, this means different interrupt handling routines depending on baud rates. To do this any other way lessens the reliability of other high-speed interrupt-driven peripherals in the system.

And remember, it's only acceptable to patch the main interrupt vector when there is no other way to service interrupts fast enough. At all other times, even in the same program, service your interrupts in other ways.

VECTORS VS. BINDING VS. ALLOCATING

There are three main ways to get into the IIgs interrupt-handling chain--by patching vectors directly, by using the ProDOS 8 or ProDOS 16 call `ALLOC_INTERRUPT`, and by using the GS/OS call `BindInt`. Each behaves differently and has advantages and disadvantages. We'll go from the highest level to the lowest in discussing them.

BINDINT--EASY TO USE, BUT NOT AS EASY TO CONTROL

BindInt's vector reference numbers (VRNs) are designed to correspond to vectors in the IIgs Interrupt Manager's chain. Comparing the list of numbers on page 265 of GS/OS Reference to the list of vectors starting on page 266 of the Apple IIgs Firmware Reference will make this more obvious.

When you call BindInt, GS/OS replaces the address in the appropriate interrupt vector with an address inside GS/OS. The routine it points to calls all the routines bound to that vector, including the one that was originally installed (usually the ROM's built-in SEC/RTL address). That is, if IRQ.VBL pointed to the Miscellaneous Tools' Heartbeat Task code before a program made four separate BindInt calls to VRN \$000C, then after those calls completed, IRQ.VBL would point to code inside GS/OS that called all four bound routines and the Miscellaneous Tools' Heartbeat Task code.

This is why each bound routine is told (through the microprocessor's carry flag) if one of the other routines has already claimed the interrupt and why preserving that status is important. BindInt is a convenient way to get code time during various kinds of interrupts, but you should note that you can't control in what order bound handlers are called.

ALLOC_INTERRUPT--OLD STYLE INTERRUPT MANAGEMENT

ALLOC_INTERRUPT and the ProDOS 8 equivalent, ALLOC_INT take the address of the routine you pass and keep it in an internal table. When an interrupt occurs, each address in the table is called in turn until one of the interrupt handlers claims it. In older days, failure by any of the installed interrupt handlers to claim the interrupt would bring the system to a crashing halt--nowadays unclaimed interrupts are ignored by both ProDOS 8 and GS/OS.

What the manuals don't tell you is that any routine installed in this way is called after the system has jumped through address \$03FE in bank zero--in other words, at the last possible chance. For any kind of timing-sensitive interrupts, these routines are not sufficient.

The table that stores these routines is of a fixed size--ProDOS 8's table holds four routines, and GS/OS's holds 16. If you try to install more handlers than that, you'll get an error from the operating system.

PATCHING VECTORS--HIGH LEVEL OF CONTROL, HIGH RISK

The lowest level at which you can get control is by directly patching the Interrupt Manager's vectors as documented in the Firmware Reference. Although this lets you get control as soon as the Interrupt Manager determines which vector to call, it also carries some compatibility risks.

Any BindInt calls with VRNs that reference a vector you patch make GS/OS take your routine's address and store it internally. This is a problem for anyone who daisy-chained into the same interrupt vector after you did--there's no good way to disconnect yourself without disconnecting everyone who patched in after you. This is Bad.

If you patch vectors directly, you have to check the vector when you're ready to remove your routine. If the vector doesn't still point to your address, someone else has patched into the vector after you and you can't remove yourself. In these cases, you have to leave a "code stub" that takes no action other than passing control along to the address that was installed when you patched in, and you have to leave that code stub at the same address as your interrupt handler. (Since you don't know who has patched the vector after you, you have no way to communicate with those programs and tell them

you're going away.)

This means your interrupt handler can't be in your main program. If it is, when GS/OS calls UserShutDown to remove your program from memory, you'll orphan one or more pointers to your interrupt handler (which doesn't exist anymore). You must allocate memory and load your interrupt handler with a different user ID than your main program so your code stub can survive when your program quits. Also note that this means repeated launchings of your program could leave lots and lots of code stubs in memory--so if you can find a way other than patching vectors directly, you're encouraged to use it.

SOFTWARE INTERRUPTS--BRK AND COP

Sometimes developers forget that BRK and COP instructions are in fact software interrupts--when the IIgs's 65816 encounters one of these instructions, it goes through the same Interrupt Manager procedures that all interrupts go through.

Among other things, this means that encountering one of these instructions inside an interrupt routine will overwrite all the system's saved information (such as registers or system state variables) with new ones, meaning you'll never be able to return from the first interrupt. This isn't too much of a problem with BRK (except when debugging interrupt routines), but a recent fad popularity for COP makes this worth mentioning.

Some developers are trying to use COP instructions for all kinds of general-purpose mechanisms, but the system is not designed to handle this. Using a COP instruction to pass control to a shell or a library routine in production-level code is not acceptable for several reasons. First, any COP instruction inside an interrupt handler will bring the system to its knees. Second, there is no arbitration for the COP vector so multiple users of it will collide. Third, although a COP instruction takes only two bytes, it takes many hundreds more cycles to execute than a JSL instruction, slowing the system down for no reason.

COP instructions are perfectly acceptable in non-production level (debugging) code, but developers should not use them as a way for different program modules to communicate. Such use is not supported and is strongly discouraged by Apple.

BEFORE WE RTI--A SUMMARY

This Note covers many issues concerning interrupts, so here's a summary. This isn't all the explanation--refer to individual topics for discussions and reasons.

- o Never disable interrupts for longer than necessary--you make life really difficult on routines that rely on high-speed interrupt capability.
- o Interrupt routines should patch in as late as possible in the interrupt process without losing data. If your interrupt source doesn't need servicing as fast as AppleTalk does, don't patch in before AppleTalk.
- o Patching the main interrupt vector at \$E10010 is only acceptable if there's no possible way to service external interrupts quickly enough (internal interrupt sources, like serial ports, should always use other vectors), and even then the vector must only be patched while necessary; if a slower interrupt source is used in the same program,

unpatch the vector.

- o Different methods of installing interrupt handlers give you different levels of control. BindInt is the overall best method, although you can't control in what order bound routines are called.
- o COP instructions are unacceptable in non-debugging code; they should never take the place of JSL instructions or other methods of inter-process communication.

Further Reference

- o Apple IIgs Firmware Reference
- o Apple IIgs Toolbox Reference, Volume 3
- o GS/OS Reference
- o ProDOS 8 Technical Reference Manual
- o Apple IIgs Technical Note #24, Apple IIgs Toolbox Reference updates

Apple IIgs
#106: ADB Addendum

Written by: Dave Lyons

May 1992

This Technical Note documents some bits in the ADB SendInfo data byte for setModes and clearModes.

SendInfo is documented in Volume 1 of Apple IIgs Toolbox Reference, but it doesn't tell you what any of the bits in the setModes/clearModes data byte are for. Well, here are the useful ones:

Bit	Value	Description
6	\$40	Shift+CapsLock=Lowercase mode
4	\$10	Keyboard buffering
3	\$08	Dual-speed keys
2	\$04	Fast space/delete keys

For example, to turn off keyboard buffering without altering the user's Battery RAM, you can do the following:

```
        pea 1                ;number of data bytes
        pushlong #modesToClear ;pointer to data byte
        pea 5                ;modeCmd = clearModes
        _SendInfo
        ...

modesToClear dc.b $10                ;bit 4 = keyboard buffering
```

Note that the user's control panel setting will become current again if they hit Command-Ctrl-Esc (the system calls the TOBRAMSETUP vector at \$E10094 to update the system from Battery RAM).

Further Reference

- o Apple IIgs Toolbox Reference, Volume 1

Apple IIgs
#107: Tool Locator Tribulations

Written by: Dave Lyons

May 1992

This Technical Note tells you what to watch out for in the Tool Locator.

SHUTDOWNTOOLS AND SYSTEM 6.0

In System 6.0, ShutDownTools inappropriately calls HideCursor even if QuickDraw II is not started. The results are unpredictable.

If your application does not use QuickDraw II but does use ShutDownTools, you may need to start up and shut down your tools manually instead.

Note that the HideCursor problem does not occur in the (unusual) case that the System 6.0 noResourceMgr bit (value \$0010) is set.

CONTENTS OF THE STARTUPTOOLS TOOL TABLE

You should not include the Tool Locator or Memory Manager in your tool table. Instead, call TLStartUp and MMStartUp before calling StartUpTools, and call MMShutDown and TLShutDown after ShutDownTools.

Since StartUpTools automatically starts the Resource Manager for you, you should not include Resource Manager in the tool table, either. Doing so has no ill effect in System Software 6.0 and later, but in System Software 5.0 through 5.0.4 you get duplicate ResourceStartUp and ResourceShutDown calls.

The order of the tool table entries does not matter. (Toolbox Reference 3, page 51-8, says "Although StartUpTools handles the order of tool startup for you...", but this is widely overlooked.)

Further Reference

- o Apple IIgs Toolbox Reference, Volume 3

Apple IIgs
#108: Finder Funkiness

Written by: Dave Lyons

May 1992

This Technical Note tells you what to watch out for in Finder 6.0.

ICON SEARCH ORDER

When the Finder looks for an icon it uses the first match it finds. When more than one icon would match, the order is important.

Some icons are built right into the Finder's resource fork--those are always searched last. Other than that, the Finder searches in device-number order (for example, icons on device number \$0001, the boot device, override icons on other devices).

On each disk, icons in Desktop files override icons in old-style icon files. Among old-style icon files in the same Icons folder, each icon file overrides subsequent (as returned by GetDirEntry) files in the same directory. Within an icons file, earlier icons override later icons.

If you create a "generic" icon that matches a broad class of files--for example, all files of a particular file type--you have to be very careful where you put that icon. A generic icon in any file's rBundle will wind up in a Desktop file, where it will override some old-style icon files (or all of them, if the Desktop file is on the boot disk).

There's really no good place for a custom generic icon. (Well, the Finder's resource fork would be a good place, but we recommend not messing with that.) A halfway-good place is in old-style icons folders, at the end, on the highest-numbered convenient device (for example, your third hard drive partition of three).

Note that the 6.0 Finder's matching order for old-style icons is more or less the reverse of what it was in previous versions.

FILENAME MATCHING AND WILDCARDS

When an icon matches by filename and has a leading wildcard, the match always fails if there are any lowercase characters in the string. For example, "*.TXT" is fine, but "*.Txt" never matches.

Also, a leading wildcard matches one or more characters, instead of (as intended) zero or more characters. "*ICONS" matches "MyIcons" and "Other.Icons", but not "Icons". You can usually work around this by omitting the character after the wildcards: "*CONS" matches all three.

These notes apply both to old-style icon files and to new matchFilename structures.

SHUT DOWN DEFAULT IS NOT CONFIGURABLE

The System 6.0 Finder Documentation shows one of the words in the rRectList(1) resource as the default choice for the Shut Down dialog. Actually, the default is not configurable, and this word in the resource should remain zero.

Utilities can customize the Finder's "Shut Down..." command by accepting the `finderSaysMItemSelected` request.

Further Reference

- o System 6.0 Documentation