




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification ⁷ : G09G 1/16</p>	<p>A2</p>	<p>(11) International Publication Number: WO 00/46781 (43) International Publication Date: 10 August 2000 (10.08.00)</p>
<p>(21) International Application Number: PCT/US00/03165 (22) International Filing Date: 4 February 2000 (04.02.00)</p> <p>(30) Priority Data: 09/246,040 5 February 1999 (05.02.99) US 09/263,612 5 March 1999 (05.03.99) US 09/369,053 4 August 1999 (04.08.99) US</p> <p>(71) Applicant (for all designated States except US): XSIDES CORPORATION [US/US]; Suite 1600, 821 Second Avenue, Seattle, WA 98104-1504 (US).</p> <p>(72) Inventors; and (75) Inventors/Applicants (for US only): NASON, D., David [US/US]; Suite 1600, 821 Second Avenue, Seattle, WA 98104-1504 (US). O'ROURKE, Thomas, C. [US/US]; Suite 1600, 821 Second Avenue, Seattle, WA 98104-1504 (US). CAMPBELL, Scott [US/US]; Suite 1600, 821 Second Avenue, Seattle, WA 98104-1504 (US). EASTON, John [US/US]; Suite 1600, 821 Second Avenue, Seattle, WA 98104-1504 (US). KAAAN, Carson [US/US]; Suite 1600, 821 Second Avenue, Seattle, WA 98104-1504 (US). BROOKS, Phillip [US/US]; Suite 1600, 821 Second Avenue, Seattle, WA 98104-1504 (US).</p>	<p>(74) Agents: BIERMAN, Ellen, M. et al.; Seed Intellectual Property Law Group PLLC, Suite 6300, 701 Fifth Avenue, Seattle, WA 98104-7092 (US).</p> <p>(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CR, CU, CZ, DE, DK, DM, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).</p> <p>Published Without international search report and to be republished upon receipt of that report.</p>	

(54) Title: ALTERNATE DISPLAY CONTENT CONTROLLER

(57) Abstract

An alternate display content controller provides a technique for controlling a video display separately from and in addition to the content displayed on the operating system monitor. Where the display is a computer monitor, the alternate display content controller interacts with the computer utility operating system and hardware drivers to control allocation of display space and create and control one or more parallel graphical user interfaces adjacent the operating system desktop. An alternate display content controller may be incorporated in either hardware or software. As software, an alternate display content controller may be an application running on the computer operating system, or may include an operating system kernel of varying complexity ranging from dependent on the utility operating system for hardware system services to a parallel system independent of the utility operating system and capable of supporting dedicated applications. The alternate display content controller may also include content and operating software delivered over the internet or any other LAN. The alternate display content controller may also be included in a television decoder/settop box to permit two or more parallel graphical user interfaces to be displayed simultaneously.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

ALTERNATE DISPLAY CONTENT CONTROLLER

BACKGROUND OF THE INVENTION

5 1. Field of the Invention

This invention relates to user interface displays and, in particular, the use of one or more parallel user interfaces separate from the standard user interface display.

2. Description of the Prior Art

10 There was a time when the most popular operating system for personal computers (DOS) did not include a graphical user interface. Any company could create a "menu" or "shell" which would be the first program launched upon starting the computer and which would present options to the user for launching and managing various applications. Although graphics programming was difficult in the DOS environment, some companies even created graphical user interfaces that could then launch other programs.

15 Microsoft Corporation of Redmond, Washington, introduced such a graphical user interface for launching applications which it called "Windows". The first three versions of Windows were merely applications which ran under DOS and could be one of numerous items to be selected from a previously running shell or menu which might be offered by a company other than Microsoft. This continued to allow other companies to offer primary user interface
20 programs to users without the user going through a Microsoft controlled user interface.

However, with the introduction by Microsoft of Windows 95™, the initial loading of the operating system presents a Microsoft-developed graphical user interface (GUI) at the outset, which occupies the entire screen display. This operating system created GUI is commonly known as a "desktop". As with its previous operating system products, Microsoft arranged with
25 manufacturers of the standard computer hardware to include this operating system with each computer sold. Microsoft's OEM licensing restrictions prevent vendors from altering, obscuring, or preceding the Microsoft desktop display. The Windows environment also presumes its ownership of the entire display and is designed in ways that assume that it can write to any

screen location at any time. With Microsoft's domination of this market, it became impossible for other software vendors to present an interface to users other than as a Microsoft style icon within the Microsoft "desktop" consisting of the entire screen display. This prompted a need for access to a user interface which could be presented outside of the standard computer screen display and therefore independent of the dictates of Microsoft for items within its "desktop".

Standard personal computers use VGA or Super VGA or XGA video display systems. These display systems operate in standardized graphics modes such as 640 x 480 pixels, 800 x 600 pixels, 1024 x 768 pixels, and 1280 x 1024 pixels. When one of these display modes is selected, this is the entire area available for display. In the Microsoft Windows environment, the user instructs the Windows operating system to select one of these standard display modes and the Windows operating system then presents all of the applications and their icons within the selected display area. There is no way at present to cause the Windows "desktop" to use less than the entire display area and still function as intended and allow another program from another vendor to control the remainder. What is needed is the ability to designate a portion of video memory a separate from the Windows desktop, and to make sure that Windows functions normally but at the same time cannot obstruct anything subsequently allocated into that space

SUMMARY OF THE INVENTION

A first aspect of the present invention includes a technique for controlling allocation and content of display space among one or more user interfaces, operating systems or applications permitting an application or parallel graphical user interface (GUI) to operate outside the desktop, the area designated for display of the operating system interface and its associated applications. In a first aspect, a computer operating under the control of any utility operating system such as Microsoft Windows™, Linux, Apple O/S or Unix may have the allocation of visible display controlled by the present invention. The operating system desktop may be scaled and/or moved to a specific area of the display permitting a parallel GUI to operate in the open area. The present invention may be an application operating under the primary or utility operating system or it may be combined with an operating system kernel to control the display and content in the parallel display.

Another aspect of the present invention includes a technique provided for adding and using a parallel graphical user interface adjacent to the standard user graphical display interface, for example in the border beyond the standard screen display area. Conventional video systems, such as VGA, SVGA and XGA video systems, include a defined border surrounding the display area. The original purpose of this border was to allow adequate time for the horizontal and vertical retrace of the electron gun in a cathode ray tube display. However, with the advent of LCD displays and as retrace speeds have increased in modern monitors, it is now possible to present a user interface display in this border. The border which can be controlled as a user interface is a portion of what is known as the "overscan". This invention is a method for presenting one or more additional or secondary user interfaces, for example, in the overscan area surrounding the conventional user interface display often called the desktop.

When the electron gun in a CRT retraces to the left of the screen or the top of the screen, it requires a significant amount of time relative to the presentation of a scanned line of data. During the retrace, the electron gun is turned off ("blanked"). If the blanking time required for the retrace is equal to the amount of time available, there is no usable overscan. However, modern monitors have become much faster in their retrace speeds, leaving a significant amount of time when the electron gun need not be blanked, allowing a displayable border. In the prior art, although the border is usually "black" (the gun is turned off), it is well known how to specify that the border shall be given any one of six colors. Standard BIOS allows a specification of this color. The desired color is simply specified in one of the registers for the video controller. Typically no data for this color is stored in the buffer of video memory for the display. This invention establishes an additional video buffer for the border and allows this buffer to be written with display data like the regular display buffer. The additional video buffer is often present but unused in the graphics systems of most computers because video memory is usually implemented in sizes that are powers of 2 e.g. "512K", whereas standard desktop dimensions are not "e.g. 640x480=300K". The display area is thereby expanded, on one or more edges, to provide a visible area previously invisible. The pixels within this newly visible area of the display are made accessible to programs through an application programming interface (API) component of this invention. A program incorporating a parallel graphical user interface may be displayed in the previously blanked area of the display, functionally increasing the accessible

area of the display without hardware modification. In other cases the desktop may be increased or decreased to non-standard sizes.

A further aspect of the present invention includes a method for displaying an image on a video display system in an area outside of the primary display area generated by the video display system. Two dimensions define the standard display area, each specifying a number of pixels. Selecting a video "mode" specifies these dimensions. The method is accomplished by adjusting parameters for the video display system to increase the number of pixels in at least one dimension of the display system. The number of pixels which is added is less than or equal to the difference between the number of pixels specified in the video mode and a maximum number of pixels which the video display system can effectively display. Any such difference is defined here as the overscan area. Thus, the overscan area may be the difference between the current desktop video mode and the display capability of the display device or more specifically, any portion of video memory unused when the operating system is in a given screen dimension. Because all interface displays are created by writing a desired image to a buffer or memory for the video display, the method requires allocating additional video display memory for the increased pixels. The image written to such memory is then displayed by the system alongside the original display area.

In a still further aspect of the present invention, only the vertical dimension is increased and the overscan user interface is presented above or below the primary display area. Alternatively, the horizontal dimension may be increased and the overscan user interface displayed to the right or the left of the primary display area. Similarly, the interface image may be displayed on any or all of the four sides of the primary display area.

In another still further aspect of the present invention, a parallel GUI is provided that includes access to existing search engines and browsers. In another embodiment, the parallel GUI includes a search engine and/or browser. A search engine and/or browser using the present invention may be opened in either the overscan space or a space within or over the operating system display.

These and other features and advantages of this invention will become further apparent from the detailed description and accompanying figures that follow. In the figures and description, numerals indicate the various features of the invention, like numerals referring to like features throughout both the drawings and the description.

5 **BRIEF DESCRIPTION OF THE DRAWINGS**

Fig. 1 is a block diagram of a first embodiment of the present invention.

Fig. 2 is a block diagram of a second embodiment of the present invention.

Fig. 3 is a diagram of a standard display with an overscan user interface on all four borders of the display.

10 Fig. 4 is a block diagram of the basic components of the present invention.

Fig. 5 is a diagram of a cursor or pointer within the overscan user interface and the hotspot above it within the standard display.

Fig. 6 is a diagram of the usable border within the vertical overscan and the horizontal overscan surrounding the standard display.

15 Fig. 7 is an overview flow chart showing the operation of a preferred embodiment of the present invention.

Fig. 8 is a flowchart of the sub-steps in Identify Display step 102 of Fig. 7.

Fig. 9 is a flowchart of the sub-steps of changing the display resolution step 114 of Fig. 7.

Fig. 10 is a flowchart of the sub-steps in the Paint the Display step 120 of Fig. 7.

20 Fig. 11 is a flowchart of the sub-steps of Enable Linear Addressing step 112 of Fig. 7.

Fig. 12 is a flowchart of the sub-steps of the Process Message Loop of Fig. 7.

Fig. 13 is a flowchart of the sub-steps of the Check Mouse and Keyboard Events step 184 in Fig. 12.

Fig. 14 is a flowchart of the sub-steps of the Change Emulation Resolution step 115 in Fig. 7.

5 Fig. 15 is a diagram of a standard display of the prior art.

Fig. 16 is a diagram of a standard display with an overscan user interface in the bottom overscan area.

Fig. 17 is a diagram of a standard display including a desktop, an overscan user interface in the bottom overscan area and a context sensitive browser on the side.

10 Fig. 18 is a diagram of a standard display with an overscan user interface in the bottom and on the right overscan area.

DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT

The present invention includes techniques for providing and using an additional user interface, preferably a secondary graphical user interface or parallel GUI, to be present on the display at least apparently simultaneously with the primary user interface, such as the conventional desktop GUI.

Referring now to Fig.'s 1 and 2, in a preferred embodiment, programming mechanisms and interfaces in a video display and control system such as computer system 7 or settop box 8 provide one or more parallel GUIs such as space 2C and/or space 4 in a display area such as display area 1 or display area 9 by providing access and visibility to a portion of the display otherwise ignored and/or inaccessible (hereinafter "overscan area"). Display areas such as display area 1 or display area 9 may be created on any type of analog or digital display hardware including but not limited to CRT, TFT, LCD and flat panel.

Alternate display content controller 6 interacts with the computer utility operating system 5B and hardware drivers 5C to control allocation of display space 1 and create and control one or more parallel graphical user interfaces such as context sensitive network browser (CSNB) 2 and internet pages 2A and 2B adjacent the operating system desktop 3. Alternate display content
5 controller 6 may be incorporated in either hardware or software. As software, an alternate display content controller may be an application running on the computer operating system, or may include an operating system kernel of varying complexity ranging from dependent on the utility operating system for hardware system services to a parallel system independent of the utility operating system and capable of supporting dedicated applications. The alternate display
10 content controller may also include content and operating software such as JAVA delivered over the Internet I or any other LAN.

The alternate display content controller may also be included in a television decoder/settop box such as box 8 to permit two or more parallel graphical user interfaces such as pages 9A and 9B to be displayed simultaneously. The present invention may be compatible with
15 conventional television formats such as NTSC, PAL, PAL-C, SECAM and MESECAM. In this configuration content and software may be delivered over any conventional delivery medium 10 including but not limited to over the air broadcast signals 10A, cable 10C, optical fiber, and satellite 10B.

Fig.'s 1 and 2 will be referenced in more detail later in the application.

20 Fig. 15 shows a standard prior art display desktop generated by a Microsoft Windows 95™ operating system. Within the desktop 31 are the taskbar 32 and desktop icons 33.

In a preferred embodiment of the present invention, a graphical user interface image is painted onto one or more of the sides of the overscan area as shown in Fig. 3. Fig. 3 is a depiction of a Super VGA (SVGA) display with the addition of a graphical bar user interface
25 displayed in the overscan area. The overscan user interface bar 30 is defined to reside outside the borders of the "desktop" display area 31. In Fig. 16, the display is modified to include a graphical user interface 30 in a bar 20-pixels high below the bottom edge. In Fig. 3, the display is modified to include a graphical user interface in four bars each 20-pixels high/wide outside

each of the four display edges: a bottom bar 30, a left side bar 34, a right side bar 36, and a top bar 38.

The overscan interface may include, and is not limited to, buttons, menus, application output controls (such as a "ticker window"), animations, and user input controls (such as edit boxes). Because the overscan interface is not obscured by other applications running within the standard desktop, the overscan interface may be constantly visible or it may toggle between visible and invisible states based upon any of a number of programming parameters (including, but not limited to, the state of the active window, the state of a toggle button, etc.).

Fig. 4 is a block diagram of the basic components of the present invention. Within the software component S are the operating system 63 and one or more applications such as application 61. Within the protected modes of modern systems, applications 61 do not have direct access to the video or Graphics Drivers 64 or hardware components such as the video card 66 which contains the video chipset 66A, 66B and 66C. Abstraction layers such as Application Interface (API) 60, and/or Direct API 62, provide limited access, often through the operating system 63.

The invention provides a technique for painting and accessing an area of the computer display not accessible, or used, in the operative desktop graphics modes. In the Microsoft Windows environments (including Microsoft Window 95 and derivatives, and Microsoft Windows NT 4.0 and derivatives) and other contemporary operating environments, the primary display area "desktop" is usually assigned by the operating system to be one of a set of predetermined video "modes" such as those laid out in Tables 1 and 2 below, each of which is predefined at a specific pixel resolution. Thus, the accessible area of the computer display may not be modified except by selecting another of the available predefined modes.

TABLE 1: ROM BIOS video modes

Mode Number	Resolution	Mode Colors	Buffer Type	Segment
00H	42x25 chars (320x350 pixels)	16	Alpha	B800

00H	42x25 chars (320x350 pixels)	16	Alpha	B800
00H	42x25 chars (320x400 pixels)	16	Alpha	B800
00H	42x25 chars (320x400 pixels)	16	Alpha	B800
01H	42x25 chars (320x200 pixels)	16	Alpha	B800
01H	42x25 chars (320x350 pixels)	16	Alpha	B800
01H	42x25 chars (320x400 pixels)	16	Alpha	B800
01H	42x25 chars (320x400 pixels)	16	Alpha	B800
02H	80x25 chars (640x200 pixels)	16	Alpha	B800
02H	80x25 chars (640x350 pixels)	16	Alpha	B800
02H	80x25 chars (640x400 pixels)	16	Alpha	B800
02H	80x25 chars (640x400 pixels)	16	Alpha	B800
03H	80x25 chars (640x200 pixels)	16	Alpha	B800
03H	80x25 chars (640x350 pixels)	16	Alpha	B800
03H	80x25 chars (640x400 pixels)	16	Alpha	B800
03H	80x25 chars (720x400 pixels)	16	Alpha	B800
04H	320x200 pixels	4	Graphics	B800
05H	320x200 pixels	4	Graphics	B800
06H	840X200 pixels	2	Graphics	B800
07H	80x25 chars (720x350 pixels)	2	Alpha	B000
07H	80x25 chars (720x400 pixels)	2	Alpha	B000
0DH	320x200 pixels	16	Graphics	A000
0EH	640x200 pixels	16	Graphics	A000

0FH	640x350 pixels	4	Graphics	A000
10H	640x350 pixels	4	Graphics	A000
10H	640x350 pixels	16	Graphics	A000
11H	640x480 pixels	2	Graphics	A000
12H	640x480 pixels	16	Graphics	A000
13H	320x200 pixels	256	Graphics	A000

TABLE 2: SVGA video modes defined in the VESA BIOS extension

Mode Number	Resolution	Mode Colors	Buffer Type
100H	640x480 pixels	256	Graphics
101H	640x480 pixels	256	Graphics
102H	800x600 pixels	16	Graphics
103H	800x600 pixels	256	Graphics
104H	1024x768 pixels	16	Graphics
105H	1024x768 pixels	256	Graphics
106H	1280x1024 pixels	16	Graphics
107H	1280x1024 pixels	256	Graphics
108H	80x60 chars	16	Alpha

109H	132x25 chars	16	Alpha
10AH	132x43 chars	16	Alpha
10BH	132x50 chars	16	Alpha
10CH	132x60 chars	16	Alpha
10DH	320x200 pixels	32,768	Graphics
10EH	320x200 pixels	65,536	Graphics
10FH	320x200 pixels	16,777,216	Graphics
110H	640x480 pixels	32,768	Graphics
111H	640x480 pixels	65,536	Graphics
112H	640x480 pixels	16,777,216	Graphics
113H	800x600 pixels	32,768	Graphics
114H	800x600 pixels	65,536	Graphics
115H	800x600 pixels	16,777,216	Graphics
116H	1024x788 pixels	32,768	Graphics
117H	1024x768 pixels	65,536	Graphics

118H	1024x768 pixels	16,777,216	Graphics
119H	1280x1024 pixels	32,768	Graphics
11AH	1280x1024 pixels	65,536	Graphics
11BH	1280x1024 pixels	16,777,216	Graphics

As shown in Fig. 6, a displayed image is "overscanned". That is, the displayed video buffer data occupies less than the entire drivable screen size. The drivable screen size is determined by the total amount of video memory and the operative video display characteristics.

5 The width of the usable overscan border depends on the amount of the horizontal overscan 52 reduced by the horizontal blanking 54 and the amount of the vertical overscan 53 reduced by the vertical blanking 55.

In a first preferred embodiment, only a border at the bottom of the standard display area is used. Consequently, only the vertical control parameters for the cathode ray tube (CRT) controller, shown as Control Registers 6H, 16H, 11H, 10H, 12H and 15H in Fig. 4 need to be adjusted. These parameters and others are shown in Table 3 below:

10

TABLE 3: Vertical timing parameters for CR programming.

Register	Name	Description
6H	Vertical Total	Value = (total number of scan lines per frame) – 2 The high-order bits of this value are stored in the overflow registers.
7H	Overflow	High-order bits from other CR registers.
10H	Vertical Retrace Start	Scan line at which vertical retrace starts. The high-order bits of this value are stored in the overflow registers.
11H	Vertical Retrace End	Only the low-order 4 bits of the actual Vertical Retrace End value are stored. (Bit 7 is set to 1 to write-protect registers 0 through

		7.)
12H	Vertical Display End	Scan line at which display on the screen ends. The high-order bits of this value are stored in the overflow registers.
15H	Start Vertical Blank	Scan line at which vertical blanking starts. The high-order bits of this value are stored in the overflow registers.
16H	End Vertical Blank	Scan line at which vertical blanking ends. The high order bits of this value are stored in the overflow registers.
59H- 5AH	Linear Address Window Position	Linear address window position in 32-bit CPU address space.

In the standard 640x480 graphics mode, the nominal horizontal scan rate is 31.5 KHz (31,500 times per second) with a vertical scan rate of 60 Hz (60 frames per second). So the number of lines in one frame is 31,500/60, or 525. Because only 480 lines of data need to be displayed, there are 525-480, or 45, lines available for vertical overscan. Leaving a more than adequate margin for retrace, which requires only 2 lines worth of time, the preferred embodiment uses 20 lines for the alternate display. Thus the additional 23 unused but available lines may be used to increase the size of the operating system desktop to some non-standard size while still allowing two lines for retrace, or may be left blank, or may be used for one or more additional alternate parallel user interface displays.

The disclosed method of the preferred embodiment of the present invention is accomplished by achieving three requirements:

- (1) to address and modify the visible resolution of the video display system such that portions of the overscan area are visible as shown in Fig. 6,
- (2) to address and modify the video display contents for the visible portion of the overscan area, and
- (3) to provide an application programming interface (API) or other mechanism to allow applications to implement this functionality.

Fig. 7, and the additional details and sub-steps provided in Fig.'s 8-13, provides a flow chart of an implementation of a preferred embodiment of the present invention meeting the requirements described above. The environment of this implementation is a standard Microsoft Windows 95™ operating environment, using Microsoft Visual C and Microsoft MASM for the development platform. That is not to imply that this invention is limited in scope to that environment or platform. The invention could be implemented within any graphical interface environment, such as X-Windows, OSF Motif, Apple OS, a Java OS, and others in which similar video standards (VGA, SVGA, XGA, 8514/A) are practiced. The reference books PC Video Systems by Richard Wilton, published by Microsoft Press and Programmer's Guide to the EGA, VGA, and Super VGA Cards by Richard F. Ferrano, published by Addison Wesley provide more than adequate background information to implement this embodiment.

Referring now in particular to Fig. 7, upon initialization, at Identify Display Type step 102, the program attempts to determine the display type, and current location in memory used by the display driver, in order to determine the size and locations of any display modifications to be made, *e.g.* to the size and location of overscan area(s) to be used.

As described in further detail in Fig. 8, the program first queries the hardware registry in Query Hardware Registry, step 131, to attempt to determine the registered display type. If successful, the program then determines compatibility information in Display Type Supported, step 135, to verify that the program supports that display type and determine memory allocation information.

If the hardware registry information is unavailable, as determined in step 131, or the display type determined in step 131 is unsupported as determined by step 104, the program may use an alternate approach, shown as subroutine Query hardware, steps 135 in Fig. 8, to query the BIOS, in step 134, and the video chipset 66, in step 136, for similar information as described immediately below.

If the BIOS is to be accessed in step 134, physical memory is first allocated in Allocate Physical Memory, step 132, and accessed using Microsoft's DPMI (DOS Protected Mode Interface) to map it to the linear memory address in which the BIOS resides in Use DPMI to assign BIOS linear address to physical memory, step 133.

Thereafter, the program queries the BIOS in Read BIOS block, Search for VGA/XVA type and manufacturer ID, step 134. If successful, the driver and chipset are then further queried to determine the display type and memory location in Query driver/chipset for exact chipset, step 136.

5 If the compatibility information does not indicate a standard VGA, SVGA, XGA, or 8514/A signature, step 134, this routine returns a failure. If a known chipset manufacturer's identification is found, the driver and/or chipset may be queried with manufacturer-specific routines, step 136, to identify and initialize, as necessary, the specific chipset.

10 If, at step 104, the program was unable to finally unable to identify the display type, either because the registry query in step 131 or the hardware query in step 135 was unsuccessful, the user may be prompted at Run in windowed mode, step 116, as to whether the program should continue to run as a standard "application bar" or "toolbar". The program may either exit or proceed to run as a toolbar on the desktop.

15 Returning now to Fig. 8, if a supported display type is detected, the program then determines the screen borders to be accessed in Identify borders to display in overscan, step 106, based upon user preferences, and determines, as necessary, whether sufficient video memory exists to make the necessary display changes. For example, if the screen is currently set to a 1024x768 resolution at 16 bits-per-pixel, and the program is to include four graphical interface bars, one on each edge, with each bar 20 pixels deep, the program must check that video memory
20 is greater than 1.7MB (required number of bytes = Pixels Width * BitsPerPixel * PixelsHeight).

The controller registers 6H, 16H, 11H, 10H, 12H and 15H as shown in Fig. 4 and detailed in Table 3, may be accessed through standard input/output ports, using standard inp/outp functions. The CR registers 6H, 16H, 11H, 10H, 12H and 15H must first be unlocked, as indicated in Unlock CRTC registers, step 108 in Fig. 7, to make them writeable. They are
25 unlocked by clearing bit 7 in controller register 11H.

Addressing of video memory, step 112, is accomplished through one of several means. One is to use the standard VGA 64 Kb "hardware window", moving it along the video memory buffer 67 (Fig. 4) in 64Kb increments as necessary. The preferred method is to enable linear

addressing by querying the video chipset for the linear window position address, step 138 of Fig. 11. This 32-bit offset in memory allows the program to map the linear memory to a physical address, steps 140 and 142 of Figure 11, that can be manipulated programmatically.

At this point the program can modify the size of the display, step 114 and Fig. 9, to
5 include the border areas. This routine first checks to determine whether or not the system is running in "toolbar" mode, step 144, and, if so, returns true. If not, it then determines whether to reset all registers and values to their original state, effectively returning the display to its original appearance, step 152. The determination is based upon a number of parameters, such as whether
10 the current resolution, step 146, reflects a standard value or previous programmatic manipulation, step 148. If a standard resolution is already set, the variables are reset to include the specified border areas, step 150. The CR registers are adjusted, step 154, to modify the scanned and blanked areas of the display. If the top or side areas are modified, existing video memory is moved accordingly in step 162 of Fig. 10.

If any of the foregoing routines returns a failure, the program may prompt the user to
15 determine whether "emulation" mode, step 13, or windowed mode step 116 should be used or if the program should exit at step 124.

In its simplest form, the invention can be treated as a technique for adding a secondary GUI by reconfiguring the actual display mode to add a modified, non-standard GUI mode in which the standard display size or resolution has been adjusted to include a secondary display in
20 addition to the primary display. For example, a standard 640x480 display is modified in accordance with the present invention to become a larger display, one section of which corresponds to the original 640x480 display while another section corresponds to a 640x25 secondary GUI display.

There are various techniques or mechanisms required for modifying the system to include
25 the secondary GUI, depending upon the requirements of the secondary GUI and upon the present circumstances of the unmodified system.

In another embodiment of the present invention system resources are allocated for a secondary GUI by fooling the video driver into going to larger resolution. This technique

automatically guarantees that enough space is kept clean, since the video driver allocates system resources according to the resolution that the video driver believes it will be operating in. To operate one or more secondary user interfaces in one or more areas of the screen it is necessary to have the memory that was associated in video memory or in the frame buffer with that location, contiguously below the primary surface free and available. By writing a series of small applets specific to hardware known to have system resource allocation problems for a secondary user interface, the secondary user interface application may run such applet whenever resolutions will be switched, initializing the chip set pertinent to that particular applet. If the application finds an applet pertinent to the current particular chip set it will be launched. The applet or minidriver initializes itself, performs the necessary changes to the driver's video resolution tables, forces a reenable, and sufficient space is subsequently available for one or more secondary user interfaces.

When reenabled, the driver allocates video memory as needed for the primary display according to the data on the UCCO resolution tables. Therefore, the modified values result in a larger allocation. Once the driver has allocated memory necessary for the primary surface, the driver will allow no outside access to the allocated memory. Thus by fooling the driver into believing that it needs to allocate sufficient memory for a resolution exactly x bytes larger than the current resolution where x is the size of one or more secondary user interfaces, the application can be sure that no internal or external use of the allocated memory location can conflict with the secondary user interface.

This method ensures that system resources will be allocated for one or more secondary user interfaces by writing an applet that would address the video driver in such a way as to force the video driver, on its next reenable, to allocate video memory sufficient for a resolution higher than the actual operating system resolution. This may also be done by modifying each instance of the advertised mode tables, and thus creating a screen size larger than the primary user interface screen size.

This technique has an additional benefit of eliminating the need to prevent the driver from actually shifting into the specified larger resolution, handing the primary user interface a larger display surface resolution. The "hardware mode table," a variant of the aforementioned

video resolution tables, is not advertised and not accessible. Therefore, when the driver validates the new resolution, checking against the hardware mode table, it will always fail and therefore refuse to shift into that resolution. Because this technique modified the advertised video resolution tables early enough in the driver's process, allocated memory was modified, and
5 memory addresses set before the failure in validate mode. Subsequently when the CRTCs are modified, in step 114, the driver is reserving sufficient memory for one or more secondary user interfaces and not making it available for any other process or purpose.

In yet another embodiment of the present invention, an enveloping driver is installed to sit above the existing driver and shims itself in between the hardware abstraction layer and the
10 actual video driver in order to be able to handle all calls to the video driver and modify the driver and the driver's tables in a much more generic fashion rather than in a chipset specific fashion. The enveloping driver shims into the primary video driver, transparently passing calls back and forth to the primary video driver. The enveloping driver finds the video resolution tables in the primary video driver which may be in a number of locations within the driver. The enveloping
15 driver modifies the tables (for example, increasing 800 by 600 to 800 by 620). A 1024 by 768 table entry may become 1024 by 800.

Like the previously described embodiment, the primary driver cannot validate the new resolution and therefore cannot actually change the display setting. As a result, the driver allocated memory, allocated the cache space, determined memory address and moved cache and
20 offscreen buffers as necessary. So the primary driver never uses all the space allocated, and will never draw in that space.

As stated earlier, the method of the present invention may include three primary steps, finding or producing unused video memory, creating or expanding the overscan area, and putting data in the overscan area.

25 The step of finding or producing the unused video memory requires a review of the contents of the Controller Registers, the CR registers, used by VGA compatible chip sets or graphic boards to identify where the overscan area, the blanking, the vertical and horizontal total and the sinking should be set. The CR defines the desktop display, how its synched, where it's

laid out left and right, how much buffer area there would be on each side, where it would be stored within the video memory area. A review of the contents of the CR data registers therefore fully defines and allows one to control the potential location and size of the overscan area.

5 In order to accomplish the step of creating or expanding the overscan area, the CRs may currently be used directly for systems with video display resolutions up to and including 1024 pixels in any dimension, that is, resolutions which can be defined in the generally accepted VGA standards by 10 bits per register. To expand the overscan area, new data is written into the CR using standard techniques such as the Inp and Outp, functions. A standard video port and MMIO functions may also be used to modify the CRs.

10 At greater resolutions, 11 bits may be needed to properly define the resolution. There is currently no standard way in which the 11th bit location is defined. Therefore, at a resolution above 1280 by 1024, for example, an understanding about the video card itself, particularly how the 11 bits representing the resolution are stored, is currently required and will be described below in greater detail.

15 When expanding the overscan, it is important to make sure a previous overscan bar is not already displayed, possibly from a previous crash or other unexpected problem. Either the display must be immediately reset to the appropriate resolution defaults, or the CR queried to determine if the total screen resolution as understood by the video card and drivers differs from the screen resolution known by the operating system display interface. An overscan bar may
20 already be displayed if the total screen resolution is not equal to one of the standard VGA or SVGA resolutions. In particular, if the total screen resolution is equal to a standard VGA/SVGA resolution plus the area required for the overscan bar or is greater than the resolution reported by the operating system display interface, the display is reset.

25 Once the display area or resolution as stored in the CR is determined, the resolution or display area can be extended in several different ways. The overscan area can be added to the bottom, the top, or the right of the current display area, and optionally, the display area can be repositioned so that the overscan bar can remain centered in appearance. Alternatively, the overscan area can be added anywhere and the original or desktop display area can be centered to

improve appearance. In any event, the height/width of the display area required for the overscan bar is presented adjacent the desktop area stored in the CR and the combination is written into the CR, overwriting the previous data.

5 The screen typically shows a quick flash as it is placed in a different mode, including the desktop display area as well as a parallel GUI such as a display bar in the overscan area. As soon as that change occurs, a black mask can be positioned over the new areas. The new menu data can then be safely written on top of the black mask so that the user never sees memory "garbage".

10 There is typically a few seconds of load time during which a simple message can be displayed, such as "Loading...", to avoid confusing the user.

There are a number of mechanisms by which this may be done. A set of class objects is used, all derived from a common base class corresponding to the above described VGA-generic technique.

15 The first mechanism is an implementation of the VGA-generic technique. Using this mechanism, no information specific to a video-card is necessary, other than ensuring VGA support. Using standard application programming interface (API) routines, primary and secondary surfaces are allocated. The new display data in the CR is simply the physical address at the start of the primary surface plus the number of pixels defined by the screen size.

20 Allocation of the primary surface will always be based on the entire screen display. Given the linear address of the allocated primary surface, from which a physical address can be derived, it can be extrapolated that the physical address of the location in video memory immediately adjacent to the primary surface is represented by the sum of the number of bytes of memory used to maintain the primary surface in memory plus the value of the physical address of the primary surface.

25 Once the physical address of the primary surface is known, the size of the primary surface as represented in video memory can be determined.

For example, the system looks in the CRs for the resolution of the screen, 800 by 600, in terms of number of bits per pixel, or bytes per pixel. Then any data stored in the CR representing any horizontal synching space is added. This is the true scan line length. The scan line length is a more accurate measurement of the width in a given resolution.

5 Next, the physical address of the allocated secondary surface is derived from its linear address. In the case where the allocated secondary surface is, in fact, allocated in the memory space contiguous to the primary surface (the value of the secondary surface physical address is equal to the value of the primary surface physical address plus the size of the primary), the secondary surface is determined to be the location in memory for the overscan display.

10 If, however, the above is not true and the secondary surface is not contiguous to the primary surface, another approach mechanism is required.

To summarize, the first mechanism determines how much physical area to allocate for the desktop allowing adjacent area for parallel GUI secondary space beyond that to display in the overscan area. The newly allocated area will be the very first block of memory available. If this
15 block immediately follows the primary surface, the physical address will correspond to the value associated with the physical address of the primary surface, plus the size of the primary surface. If that is true, the memory blocks are contiguous, this VGA-generic mechanism can be used.

If this first, VGA-generic mechanism cannot be used, the video card and driver name and version information retrieved from the hardware registry and BIOS, as described earlier, is used
20 in conjunction with a look-up table to determine the best alternatives among the remaining mechanisms. The table includes a set of standards keyed to the list of driver names found in the hardware registry. A class object specific to the video chipset is instantiated based, directly or indirectly, on the VGA-generic object.

If the hardware look up does not result in a reliable match, a reliability, or confidence,
25 fudge factor may be used. For example, if the hardware look up determines that an XYZ-brand device of some kind is being used, but the particular XYZ device named is not found in the look up table, a generic model from that chipset manufacturer may often be usable. If no

information is available, the user may get a message indicating that the hardware is not supported and that the program cannot run in the overscan area. The user may then be queried to determine if the system should be operated in the "application-toolbar" mode, which basically runs with exactly the same functionality but in a windowed environment within the desktop rather than in
5 the overscan area outside the desktop.

The next alternative mechanism uses surface overlays. The first step to this approach is to determine if the system will support surface overlays. A call is made to the video driver to determine what features are supported and what other factors are required. If surface overlays are supported, for example, there may be a scaling factor required.

10 For example, a particular video card in a given machine, using 2 megabytes of video RAM, might support unscaled surface overlays at 1024x768 at 8 bits per pixel, but not at 1024x768 at 16 bits per pixel because the bandwidth of the video card or the speed of the card, coupled with the relatively small amount of video memory would not be sufficient to draw a full width overlay. It is often horizontal scaling that is at issue, preventing the driver from drawing a
15 full width overlay. An overlay is literally an image that is drawn on top of the primary surface. It is not a secondary surface, which is described above. Typically, the system sends its signal from the video driver to the hardware such that it merges the two signals together, overlaying a second signal on top of the first.

If a system can not support unscaled overlays, perhaps because of bandwidth issues or
20 memory issues, this mechanism is not desirable. It is not rejected, but becomes a lower priority alternative. For example, if the scaling factor is below .1, then the normal bar can be drawn and it will be clipped closer to the edge. If the scaling factor is more than 10%, another approach mechanism is required

In the next set of alternative mechanisms, a secondary surface is allocated sufficient in
25 size to encompass the normal desktop display area plus the overscan area to be used for display of the overscan bar or bars. Using these mechanisms, the allocated secondary surface does not have to be located contiguous in memory to the primary surface. However, these approaches use more video memory than the others.

The first step is to allocate a secondary surface sufficient in size to contain the video display (the primary surface) plus the overscan area to be used. If the allocation fails, that means that there is not enough video memory to accomplish the task and this set of mechanisms is skipped and the next alternative tried. After the new block of memory is allocated, a timer of
5 very small granularity is used to execute a simple memory copy of in the contents of the primary surface onto the appropriate location of this secondary surface. The timer executes the copy at approximately 85 times per second.

Within this set of alternative mechanisms is a variant that uses the system page tables. This mechanism queries the system page tables to determine the current GDI surface address,
10 that is, the physical address in the page table for the primary surface. A secondary surface is then created large enough to hold all of what is in the video memory plus the memory required for the overscan bar to be displayed. This surface address is then pushed into the system page table and asserted as the GDI surface address.

Thereafter, when GDI reads from or writes to the primary surface through the driver, it
15 actually reads from or writes the new, larger surface. The overscan bar program can, subsequently, modify the area of the surface not addressed by GDI. The original primary surface can be de-allocated and the memory usage reclaimed. This mechanism, being more memory-efficient than the previously described mechanism, is the preferred alternative. But the page tables solution will not work correctly on a chipset that includes a coprocessor device. If the
20 initial device query reveals that the device does include a coprocessor, this variant mechanism will not be attempted.

Other variations of the above-described mechanisms are accounted for in derived class objects. For example, the VGA-generic mechanisms may vary when the video card requires more than ten bits to represent the video resolution in the CR. Some instances may require 11
25 bits. Such registers typically do not use contiguous bytes, but use extension bits to designate the address information for the higher order bits.

In this example, the eleventh bit is usually specified in an extended CR register and the extended CR registers are usually chip specific.

Similarly, a variation of the surface overlay mechanism includes a scaling factor, as described above. This alternative is handled in specific implementations through derived class objects and may be the best solution in certain situations.

Another implementation of this technology uses a “hooking” mechanism as shown in Fig. 14. After the display driver is identified through the hardware registry or the BIOS, as described above, certain programming interface entry points into the driver are hooked such as at step 117. In other words, when the video system device interface, Windows GDI for example, calls those entry points into the display driver, the program can take the opportunity to modify the parameters being passed to the display driver, and/or to modify the values being returned from the display driver.

By hooking the “ReEnable” function in the display driver, at step 117, the overscan bar program can allocate screen area in different ways in step 119:

- (1) In step-up mode, step 121, by intercepting a resolution change request and identifying the next-higher supported screen resolution and passing that higher resolution to the display driver, then, when the display driver acknowledges the change, intercepting the returned value, which would reflect the new resolution, and actually returning the original requested resolution instead. For example, GDI requests a change from 640x480 resolution to 800x600 resolution; the overscan program intercepts the request and modifies it to change the display driver to the next supported resolution higher than 800x600, say 1024x768. The display driver will change the screen resolution to 1024x768 and return that new resolution. The overscan program intercepts the return and instead passes the original request, 800x600, to GDI. The display driver has allocated and displays a 1024x768 area of memory. GDI and Windows will display the desktop in an 800x600 area of that display, leaving areas on the right and bottom edges of the screen available to the overscan program.
- (2) In shared mode, step 123, by intercepting only the return from the display driver and modifying the value to change the operating system’s understanding of the actual screen resolution. For example, GDI requests a change from 800x600

resolution to 1024x768 resolution. The overscan program intercepts the returned acknowledgment, subtracting 32 before passing the return on to GDI. The display driver has allocated and displays a 1024x768 area of memory. GDI and Windows will display the desktop in an 1024x736 area of that display, leaving an area on the bottom edge of the screen available to the overscan bar program.

After hooking, the overscan bar program can display by:

- (1) using standard API calls to render the bar to an off-screen buffer, as described in the next section, and then hooking the "BitBlt" function entry point into the display driver in order to modify the offset and size parameters and subsequently redirect the BitBlt to the area outside of that which the API believes is onscreen.
- (2) using mechanisms of primary and secondary surface addresses, described earlier, the program determines the linear addresses for the off-desktop memory location(s) left available to it, and can render directly to those memory locations.

Phase 2 of the invention begins by painting the new images into a standard off-screen buffer, step 118, as is commonly used in the art, and making the contents visible, step 120, as described in Fig. 10. If the program is in "toolbar" mode, step 156, the off-screen buffer is painted into the standard window client space, step 166, and made visible, step 164, using generic windowing-system routines. Otherwise, the linear window position address is mapped, step 158, as described in Fig. 11 which has been previously explained. Once the linear memory is mapped to a physical memory address, step 142, the contents of the off-screen display buffer can be copied into the video buffer directly, step 154 of Fig. 10, or painted as to a secondary surface.

The preferred embodiment application includes a standard application message loop, step 122, which processes system and user events. An example of a minimum functionality process loop is in Fig. 12. Here the application handles a minimal set of system events, such as painting requests, step 170, system resolution changes, step 172, and activation/deactivation, step 174. Here, too, is where user events, such as key or mouse events, may be handled, step 184, detailed in Fig. 13. System paint messages are handled by painting as appropriate into the off-screen

buffer, step 178, and painting the window or display buffer, step 180, as appropriate, as described earlier in Fig. 10. System resolution messages are received whenever the system or user changes the screen or color resolution. The programs reset all registers to the correct new values, then change the display resolution, step 182, as earlier described in Fig. 9, to reflect the
5 new resolution modified. User messages are ignored when the program is not the active application.

Fig. 13 describes a method of implementing user-input events. In this embodiment, there are three alternative mechanisms used to implement cursor or mouse support so that the user has a pointing device input tool within the overscan area user interface.

10 In the preferred mechanism, GDI's "cliprect" is modified to encompass the overscan bar's display area. That keeps the operating system from clipping the cursor as it moves into the overscan area. This change doesn't necessarily make the cursor visible or provide event feedback to the application, but is the first step.

Some current Windows applications continually reset the cliprect. It is a standard
15 programming procedure to reset the cliprect after use or loss of input focus. Some applications use the cliprect to constrain the mouse to a specific area as may be required by the active application. Whenever the overscan display bar interface receives the input focus it reasserts the cliprect, making it large enough for the mouse to travel down into the overscan space.

Once the cliprect has been expanded, the mouse can generate messages to the operating
20 system reflecting motion within the expansion area. GDI does not draw the cursor outside what it understands to be its resolution, however, and does not pass "out-of-bounds" event messages on to an application. The overscan program uses a VxD device driver, and related callback function, to make hardware driver calls at ring zero to monitor the actual physical deltas, or changes, in the mouse position and state. Every mouse position or state change is returned as an
25 event to the program which can graphically represent the position within the menu display bar.

An alternative mechanism avoids the need to expand the cliprect in order to avoid conflict with a class of device drivers that use the cliprect to facilitate virtual display panning.

Querying the mouse input device directly the overscan program can determine "delta's", changes in position and state. Whenever the cursor touches the very last row or column of pixels on the standard display, it is constrained there by setting the cliprect to a rectangle comprised of only that last row or column. A "virtual" cursor position is derived from the deltas available from the
5 input device. The actual cursor is hidden and a virtual cursor representation is explicitly displayed at the virtual coordinates to provide accurate feedback to the user. If the virtual coordinates move back onto the desktop from the overscan area, the cliprect is cleared, the virtual representation removed, and the actual cursor restored onto the screen.

A third alternative mechanism creates a transparent window that overlaps the actual
10 Windows desktop display area by a predefined number of pixels, for example, two or four pixels. If the mouse enters that small, transparent area, the program hides the cursor. A cursor image is then displayed within the overscan bar area, at the same X-coordinate but at a Y-coordinate correspondingly offset into the overscan area. If a two-pixel overlap area is used, this method uses a granularity of two. Accordingly, this API-only approach provides only limited vertical
15 granularity. This alternative mechanism assures that all implementations will have some degree of mouse-input support, even when cliprect and input device driver solutions fail.

Fig. 7 describes the cleanup mechanisms executed when the program is closed, step 124. The display is reset to the original resolution, step 126, and the CR registers are reset to their original values, step 128, and locked, step 130.

20 In another embodiment of the present invention, the launching or initiating of alternate display content controller 6 may be modified and controlled. Alternate display content controller 6 may be launched as a service, as an application, or as a user application. As a service, alternate display content controller 6 may be launched as a service within the registry of utility operating system 5B. The first kind of application is launched in the Run section in the registry, and the
25 user application may be initiated from the Start Up Group within the Start button. Thus, alternate display content controller 6 may be initiated any time from the first thing after graphics mode is enabled to the very last thing initiated.

Launched as a service, alternate display content controller 6 may be visible shortly after utility operating system 5B such as Windows actually addresses the display, and how soon after

depends on where alternate display content controller 6 is put it in the order of the things that will be launched as services. It may be possible to put alternate display content controller 6 so that it launches as essentially the first service and thus would launch almost at the same time as the drivers, very, very shortly after the drivers are launched. Accordingly, it is possible to have the screen change from text mode to graphics, draw the colored background, immediately re-
5 display with the overscan addressed and a parallel GUI such as CSNB 2 display the very close to the same time as taskbar. Launched as a run-line application, alternate display content controller 6 may be visible in display space 1 shortly after icons appear.

10 NetSpace

Referring again to Fig. 1, in an alternate embodiment of the present invention, the technique of controlling the allocation of display area 1 is used to open a context sensitive network browser 2 (CSNB) adjacent but not interfering with operating system desktop 3 and/or parallel graphical user interface 4. A display controller such as alternate display content
15 controller 6 may include CSNB 2 thus permitting the browser to create and control a space for itself on display 1 which may not be overwritten by utility operating system 5B. The combined controller/browser may be an application running on the computer operating system, or may include an operating system kernel of varying complexity ranging from dependent on the utility operating system for hardware system services to a parallel system independent of the utility
20 operating system and capable of supporting dedicated applications. The alternate display content controller/browser may also include content and operating software such as JAVA delivered over the Internet I or any other LAN. There may also be more than one context sensitive network browser and more than one parallel graphical user interface in addition to the operating system desktop.

25 Context sensitive interface such as network browser 2 may respond to movement and placement of cursor 1C controlled by a pointing device such as mouse 1M anywhere on display area 1. The generation and control of a cursor across two or more parallel graphical user interfaces was described previously. The location of cursor 1C will trigger CSNB 2 to retrieve appropriate and related network pages such as web page 2A. CSNB 2 may store the last X

number of CSNB enabled network addresses for display offline. In a currently preferred embodiment of the present invention, X is ten pages. If a user is examining a saved CSNB enabled page offline, a mouse click on the page or a link on the page will initiate the users dial-up sequence and establish an online connection.

5 In an alternate embodiment, alternate display content controller 6 may include a browser or search engine. In an alternate embodiment of the present invention, space 2C may include an edit input box 2D. Edit input box 2D may include conventional functionality's such as edit, copy, paste, etc.. A user may enter a URL into edit input box 2D using any conventional input device and then select a button to launch or initiate alternate display content controller 6 as a
10 browser. This may be accomplished by using objects and or drivers from utility operating system 5B. Initiating alternate display content controller 6 as a browser would include a simple window to display the URL as a live HTML document with all conventional functionality. By implementing alternate display content controller 6 as a little applet that uses that DLL, it may slide on, or slide off. Thus initiating alternate display content controller 6 as a browser is like a
15 window into the Internet.

Secondly, a user may enter any text into edit input box 2D using any conventional input device and then select a button to launch or initiate alternate display content controller 6 as a search engine. By entering a search string and selecting "search" and enter any string and click on "search" and pass that to any number from one to whatever or existing search engines, and
20 subsequently have the search string acted on by one or more selected search engines and or by alternate display content controller 6 as a search engine. Resulting in multiple different windows appearing in some sort of stacked or cascaded or tiled format, with the different searches within them.

Using alternate display content controller 6 as a search engine or browser, the results or
25 HTML document may be displayed in any overscan area or on the desktop.

Referring now to Fig 17, a context sensitive network browser such as CSNB 13 may also include a suite of tools such as tools 14 that may or may not have fixed locations on the browser space. Such tools may include but are not limited to e-mail, chat, buddy lists and voice. As

shown, spaces such as desktop 14A, web page 14B, secondary GUI 14C and browser 13 may be arranged in any convenient manner.

5 The following describes the hooking mechanism used with xSides on a Intel 80386 (or greater) processor. This description of the Intel processor operations are simplified for clarity. This hooking mechanism is expected to work on most if not all compatible processors currently available.

Interrupt Descriptor Table

10 The interrupt descriptor table (IDT) associates each interrupt with a descriptor for the instructions that service the associated event. For example, when a software interrupt (INT 3) is generated (and interrupts are enabled), the Intel processor will suspend what it was currently doing, look up in the IDT for the appropriate entry (or interrupt vector) for the address of the code to execute to service this interrupt. The code is known as the Interrupt Service Routine (ISR). It will start executing the ISR. When a Return From Interrupt instruction (IRET) is
15 executed by the ISR, the processor will resume what is was doing prior to the interrupt.

Debug Registers

The Intel 80386 microprocessor provides a set of system registers that are normally used for debugging purposes. The are technically referred to as Debug Registers. These registers allow control over execution of code as well as access over data. The Debug Registers are used
20 in conjunction with exception code. There are four addresses registers (i. e. Four different locations of code and/or data) (DR0, DR1, DR2, and DR3).

There is a control register (DR7) that can be programmed to selectively enable the address registers. In addition, DR7 is used to control the type of access to a memory location that will generate an interrupt. For example, an exception can be raised for reading and or
25 writing a specific memory location or executing a memory location (i. e. Code execution).

Finally, there is a status register (DR6) that is used to detect and determine the debug exception, (i. e. What address register generated the exception). When enabled and the data criterion is met, the x86 processor generates an Interrupt 1 (INT 1).

How this mechanism is used

5 The xSides implementation must first set up the IDT to point our ISR to process INT 1 interrupts. Next, the address of the code that you want to hook (or the memory location of data, as in this case) is programmed into one of the address registers and the appropriate bits within the control register are set. When the x86 processor executes this instruction (or touches the memory location of data), the processor generates an INT 1. The processor will then invoke the
10 Interrupt 1 ISR (as described above.) At this point, the ISR can do almost any kind of processor, code or data manipulation. When complete, the ISR executes an IRET instruction and the processor starts execution after the point of the INT 1 occurrence. Note that the interrupt code has no knowledge of the interruption.

15 This mechanism is expected to move the memory address used on some video systems for cache or hardware cursor. This should allow us to push the percentage of systems that support "overscan" mode to around 90% (in that this mechanism should work on approximately that number of machines).

Alternative embodiments

1. Utilizing the VESA BIOS Extensions (VBE) in place of the CRT Controller registers (Fig.
20 5) to determine the linear window position address, step 138, as necessary.
2. Utilizing API's (application programming interfaces) 62 capable of direct driver and/or hardware manipulation, such as Microsoft's DirectX and/or DirectDraw, in place of the CRT Controller registers and/or direct access to the display buffer.
3. Utilizing API's (applications programming interfaces) 62, such as Microsoft's DirectX
25 and/or DirectDraw, capable of direct driver and/or hardware manipulation, to create a second virtual display surface on the primary display with the same purpose, to display a separate and unobscured graphical user interface.

4. Utilizing modifications in the video subsystem of the operating system 63 in place of the CRT Controller registers and/or DirectX access to the display buffer.
5. Utilizing modifications in the video subsystem of the operating system 63 to create a second virtual display surface on the primary display with the same purpose, to display a separate and unobscured graphical user interface.
6. Building this functionality into the actual video drivers 64 and/or mini-drivers. Microsoft Windows provides support for virtual device drivers, VxDs, which could also directly interface with the hardware and drivers. These could also include an API to provide applications with an interface to the modified display.
7. Incorporating the same functionality, with or without the VGA registers, into the BIOS and providing an API to allow applications an interface to the modified display.
8. Incorporating the same functionality into hardware devices, such as monitor itself, with hardware and/or software interfaces to the CPU.
9. This technique may be used to control the desktop (i.e. Windows) to easily enable the desktop to operate in virtually any non-standard size limited only by the capability of the display hardware. This may be in combination with parallel graphical user interface displays or exclusively to maximize the primary operating system desktop display area. This may not require any modification to the operating system.

In overview, the visual display area is conventionally defined by the values maintained in the CRTC registers on the chip and available to the driver. The normally displayed area is defined by VGA standards, and subsequently by SVGA standards, to be a preset number of modes, each mode including a particular display resolution which specifies the area of the display in which the desktop can be displayed.

The desktop can only be displayed in this area because Windows does not directly read/write the video memory, rather it uses programming interface calls to the video driver. And the video driver simply reads/writes using an address that happens to be in video memory. So the value this mechanism needs to realize is the value the video card and driver assert is available

for painting. This value is queried from the registers, modified by specific amounts and rewritten to the card. Subsequently, the present invention changes the area of writable visible display space without informing the operating system's display interface of the change

5 This invention doesn't necessary change the CRTCs to add just to the bottom. Preferably the top is also moved up a little. This keeps the displayed interfaces centered within the drivable display area. For example, rather than just add thirty-two scan lines to the bottom, the top of the display area is moved up by sixteen lines.

10 Nor does this invention depend solely upon the ability to change the CRTCs to modify the visible display area. Alternative mechanisms define other methods of creating and accessing visible areas of the screen that are outside the dimensions of the desktop accessed by the operating system's display interface.

From a consideration of the specifications, drawings, and claims, other embodiments and variations of the invention will be apparent to one skilled in the art of computer science.

15 In particular, the secondary GUI may be positioned in areas not normally considered the conventional overscan area. For example, the secondary GUI may be positioned in a small square exactly in the center of the normal display in order to provide a service required by the particular system and application. In fact, the techniques of reading and rewriting screen display information can be used within the scope of the invention to maintain the primary GUI information, or portions of it, in an additional memory and selectively on a timed, computed,
20 interactive, or any or other basis, replace a portion of the primary GUI with the secondary GUI such as a pop-up, window, or any other display space.

As a simple example, a security system may require the ability to display information to a user without regard to the status of the computer system and/or require the user to make a selection, such as call for help by clicking on "911?". The present invention could provide a
25 video display buffer in which a portion of the primary GUI interface was continuously recorded and displayed in a secondary GUI for example in the center of the screen. Under non-emergency

conditions, the secondary GUI would then be effectively invisible in that the User would not notice anything except the primary GUI.

5 Under the appropriate emergency conditions, an alarm monitor could cause the secondary GUI to present the "911?" to the user by overwriting the copy of the primary display stored in the secondary GUI memory. Alternatively, a database of photographs may be stored and one recalled in response to an incoming phone call in which caller ID identified a phone number associated with a database photo entry.

10 In general, the present invention may provide one or more secondary user interfaces which may be useful whenever it is more convenient or desirable to control a portion of the total display, either outside the primary display in an unused area such as overscan or even in a portion of the primary GUI directly or by time division multiplexing, directly by communication with the video memory, or by bypassing at least a portion of the video memory to create a new video memory. In other words, the present invention may provide one or more secondary user interfaces outside of the control of the system, such as the operating system, which controls the
15 primary GUI.

Additional user interfaces may be used for a variety of different purposes. For example, a secondary user interface may be used to provide simultaneous access to the Internet, full motion video, and a conference channel. A secondary user interface may be dedicated to a local network or multiple secondary user interfaces may provide simultaneous access and data for one
20 or more networks to which a particular computer may be connected.

Having now described the invention in accordance with the requirements of the patent statutes, those skilled in this art will understand how to make changes and modifications in the present invention to meet their specific requirements or conditions. Such changes and modifications may be made without departing from the scope and spirit of the invention.

CLAIMS

- 1 1. A method for enabling the display of an image on a
2 display area of a television video display system in conjunction with the display
3 of a separately controlled user interface of the television that occupies a portion
4 of the display area, the video display system having a video display memory,
5 comprising:
6 increasing the displayable area of the video display system to
7 include a new display area portion by adjusting the parameters of the video
8 display system, thereby enlarging the display area;
9 locating additional video display memory to correspond to the
10 new display area portion, thereby creating an increased amount of video
11 display memory that is addressable;
12 allocating the enlarged display area between the separately
13 controlled television user interface and the image;
14 writing the image to the increased video display memory in
15 accordance with the allocation of the enlarged display area; and
16 transferring the video display memory contents to the video
17 display system so that the separately controlled user interface is displayed in
18 conjunction with the image.
- 1 2. The method of claim 1 wherein the enabling of the display
2 of the image is performed by computer software located in a television settop
3 box.
- 1 3. The method of claim 2 wherein the computer software is a
2 portion of the resident operating system of the television settop box.
- 1 4. The method of claim 2 wherein the computer software
2 includes a portion of an operating system separate from the resident operating
3 system of the television settop box.

1 5. The method of claim 1 wherein the enabling of the display
2 of the image is performed by computer hardware in conjunction with a
3 television settop box.

1 6. The method of claim 1 wherein the allocation of the
2 enlarged display area decreases the total area available to the separately
3 controlled user interface.

1 7. The method of claim 1 wherein the allocation of the
2 enlarged display area allocates a part of the new display area portion to the
3 separately controlled user interface.

1 8. The method of claim 1 wherein the image is displayed on a
2 part of the display area portion allocated to the separately controlled user
3 interface so that at least a portion of the image appears to continuously
4 overwrite the user interface display while the image is displayed.

1 9. The method of claim 1, the television having a
2 communications channel to an information source, the method further
3 comprising:
4 receiving data from the information source; and
5 writing the received data as part of writing the image to the
6 increased video display memory.

1 10. The method of claim 9, further comprising translating the
2 received data into a conventional television format before writing the received
3 data.

1 11. A system for enabling the display of an image on a display
2 area of a television video display system having memory, the television having
3 a separately controlled user interface that occupies a portion of the display area,
4 comprising:
5 display controller that
6 increases the displayable area of the video display system
7 to include a new display area portion by adjusting the parameters of the video
8 display system, thereby enlarging the display area;
9 locates additional memory to correspond to the new
10 display area portion, thereby creating an increased amount of memory that is
11 addressable;
12 allocates the enlarged display area between the separately
13 controlled television user interface and the image;
14 writes the image to the increased memory in accordance
15 with the allocation of the enlarged display area; and
16 transfers the increased memory contents to the video
17 display system so that the separately controlled user interface is displayed in
18 conjunction with the image.

1 12. The system of claim 11 wherein the display controller is
2 computer software located in a television settop box.

1 13. The system of claim 12 wherein the display controller
2 computer software is a portion of the resident operating system of the television
3 settop box.

1 14. The system of claim 12 wherein the display controller
2 computer software includes a portion of an operating system separate from the
3 resident operating system of the television settop box.

1 15. The system of claim 11 wherein the display controller is
2 computer hardware that performs in conjunction with a television settop box.

1 16. The system of claim 11 wherein, when the display
2 controller allocates the enlarged display area, it decreases the total area
3 available to the separately controlled user interface.

1 17. The system of claim 11 wherein, when the display
2 controller allocates the enlarged display area, it allocates a part of the new
3 display area portion to the separately controlled user interface.

1 18. The system of claim 11 wherein the image is displayed on
2 a part of the display area portion allocated to the separately controlled user
3 interface so that at least a portion of the image appears to continuously
4 overwrite the user interface display while the image is displayed.

1 19. The system of claim 11, the television having a
2 communications channel to an information source, wherein the display
3 controller, in response to receiving data from the information source, writes the
4 received data as part of writing the image to the increased memory.

1 20. The system of claim 19, wherein the display controller
2 translates the received data into a conventional television format before writing
3 the received data.

1 21. A method for enabling the display of an image on a video
2 display system in an area outside of a display area controlled by a computer
3 operating system, the computer operating system presenting a user interface
4 that occupies a portion of the display area, comprising:

5 increasing the displayable area of the video display system to
6 include a new display area portion by adjusting the parameters of the video

7 display system to increase the number of displayable pixels in at least one
8 dimension of said display area, thereby enlarging the display area;
9 locating additional video display memory to correspond to the
10 new display area portion, thereby creating an increased amount of video
11 display memory that is addressable;
12 allocating the enlarged display area between the operating system
13 user interface and the image;
14 writing the image to the increased video display memory in
15 accordance with the allocation of the enlarged display area; and
16 transferring the increased video display memory contents to the
17 video display system so that the image is displayed in conjunction with the
18 operating system user interface, wherein the method is performed by one of a
19 service of the operating system and an application program.

1 22. The method of claim 21, wherein, when the method is
2 performed as a service of the operating system, the image appears to be
3 displayed as part of the operating system initialization of the video display
4 system.

1 23. The method of claim 21, wherein, when the method is
2 performed by launching an application program, the image appears on the
3 enlarged display area when the application program is launched.

1 24. The method of claim 21, wherein the adjusting of
2 parameters of the video display system is performed by function calls to the
3 driver software for the video display system.

1 25. The method of claim 21, further comprising allocating the
2 enlarged display area between the computer operating system user interface and
3 the image, such that the portion of the display area occupied by the user
4 interface is decreased.

1 26. The method of claim 21, wherein the image is displayed
2 on a part of the display area portion allocated to the user interface so that at
3 least a part of the image appears to continuously overwrite the user interface
4 display while the image is displayed.

1 27. The method of claim 21 wherein at least a portion of the
2 image is written to the additional video memory such that the portion of the
3 image is displayed along with the operating system user interface in a manner
4 that prohibits the operating system user interface from overwriting the portion
5 of the image.

1 28. The method of claim 21 wherein the increasing of the
2 displayable area of the video display system to include a new display area
3 portion by adjusting the parameters of the video display system operates by
4 changing the resolution of the video display system and reducing the resolution
5 available to the operating system user interface.

1 29. A system for enabling the display of an image on a video
2 display system in an area outside of a display area controlled by a computer
3 operating system, the video display system having memory, the computer
4 operating system presenting a user interface that occupies a portion of the
5 display area, comprising:

6 display controller having
7 display adjustment facility that increases the displayable
8 area of the video display system to include a new display area portion by
9 adjusting the parameters of the video display system to increase the number of
10 displayable pixels in at least one dimension of said display area, thereby
11 enlarging the display area;
12 memory locator that locates additional memory to
13 correspond to the new display area portion, thereby creating an increased
14 amount of memory that is addressable;

15 display area allocation facility that allocates the enlarged
16 display area between the operating system user interface and the image; and
17 display transfer mechanism that writes the image to the
18 increased memory in accordance with the allocation of the enlarged display
19 area and transfers the increased memory contents to the video display system
20 so that the image is displayed in conjunction with the operating system user
21 interface;

22 wherein the display controller is invoked as one of a service of
23 the operating system and an application program.

1 30. The system of claim 29, wherein, when the display
2 controller is invoked as a service of the operating system, the image appears to
3 be displayed as part of the operating system initialization of the video display
4 system.

1 31. The system of claim 29, wherein, when the display
2 controller is invoked by launching an application program, the image appears
3 on the enlarged display area when the application program is launched.

1 32. The system of claim 29, wherein the display adjustment
2 facility adjusts the parameters of the video display system by performing
3 function calls to the driver software for the video display system.

1 33. The system of claim 29 wherein the display area allocation
2 facility allocates the enlarged display area such that the portion of the display
3 area occupied by the user interface is decreased.

1 34. The system of claim 29, wherein the image is displayed on
2 a part of the display area portion allocated to the user interface so that at least a
3 part of the image appears to continuously overwrite the user interface display
4 while the image is displayed.

1 35. The system of claim 29 wherein the display transfer
2 mechanism writes at least a portion of the image to the additional video
3 memory such that the portion of the image is displayed along with the operating
4 system user interface in a manner that prohibits the operating system user
5 interface from overwriting the portion of the image.

1 36. The system of claim 29 wherein the display adjustment
2 facility increases the displayable area of the video display system to include a
3 new display area portion by changing the resolution of the video display system
4 and reducing the resolution available to the operating system user interface.

1 37. A method in a computer system for enlarging a display
2 area of a video display system, comprising:
3 locating additional video display memory to correspond to a new
4 display area portion;
5 determining whether the located memory is associated with
6 another video display system function;
7 when it is determined that the located memory is associated with
8 the another video display system function, moving the use of the located
9 memory by modifying an interrupt descriptor table to capture attempts to access
10 the located memory by the another video display system function and to
11 substitute a different portion of memory for use by the another function so that
12 the located memory corresponds to the new display area portion.

1 38. A computer system for enlarging a display area of a video
2 display system having memory, the computer system having an interrupt
3 descriptor table, comprising:
4 display controller that
5 locates additional memory to correspond to a new display
6 area portion;

7 determines whether the located memory is associated with
8 another video display system function; and
9 when it is determined that the located memory is
10 associated with the another video display system function, moves the use of the
11 located memory by modifying the interrupt descriptor table to capture attempts
12 to access the located memory by the another video display system function and
13 to substitute a different portion of memory for use by the another function so
14 that the located memory corresponds to the new display area portion.

1 39. A method in a computer system for enabling the display of
2 data on a video display system in an area outside of a display area controlled by
3 a computer operating system, the computer operating system presenting a user
4 interface that occupies a portion of the display area, the data generated from a
5 network associated with the computer system, comprising:

6 increasing the displayable area of the video display system to
7 include a new display area portion by adjusting the parameters of the video
8 display system to increase the number of displayable pixels in at least one
9 dimension of said display area, thereby enlarging the display area;

10 allocating the enlarged display area between the operating system
11 user interface and the network generated data;

12 locating additional video display memory to correspond to the
13 new display area portion, thereby creating an increased amount of video
14 display memory that is addressable;

15 writing the data generated from the network to the increased
16 video display memory allocated to the network generated data; and

17 transferring the increased video display memory contents to the
18 video display system so that the data is displayed along with the operating
19 system user interface.

1 40. The method of claim 39 wherein an online connection to
2 the network is reestablished automatically when the displayed data is accessed
3 in a manner that indicates that the data is to be updated from the network.

1 41. The method of claim 39 wherein the displayed data is a
2 web page from the Internet.

1 42. The method of claim 39 wherein the displayed network
2 generated data includes an input field associated with a search engine, further
3 comprising:

4 receiving a network search request that has been entered into the
5 input field;

6 sending the received search request to the search engine;

7 receiving search results from the search engine; and

8 displaying the search results on the display area allocated for the
9 network generated data.

1 43. The method of claim 39 wherein the displayed data
2 includes an input field further comprising:

3 receiving an address of additional network data entered into the
4 input field;

5 sending the received address over the network to retrieve the
6 additional data;

7 receiving the additional data; and

8 displaying the received data on the display area allocated for the
9 network generated data.

1 44. The method of claim 39 wherein at least a portion of the
2 network generated data is written to the additional video memory such that the
3 portion of data is displayed along with the operating system user interface in a
4 manner that prohibits the operating system user interface from overwriting the
5 portion of data.

1 45. The method of claim 39 wherein the increasing of the
2 displayable area of the video display system to include a new display area

3 portion by adjusting the parameters of the video display system operates by
4 changing the resolution of the video display system and reducing the resolution
5 available to the operating system user interface.

1 46. A computer system for enabling the display of data on a
2 video display system in an area outside of a display area controlled by a
3 computer operating system, the video display system having an associated
4 memory, the computer operating system presenting a user interface that
5 occupies a portion of the display area, the data generated from a network
6 associated with the computer system, comprising:

7 display resolution adjuster that increases the displayable area of
8 the video display system to include a new display area portion by adjusting the
9 parameters of the video display system to increase the number of displayable
10 pixels in at least one dimension of said display area, thereby enlarging the
11 display area;

12 display allocation facility that allocates the enlarged display area
13 between the operating system user interface and the network generated data;

14 memory locator that locates additional memory to correspond to
15 the new display area portion, thereby creating an increased amount of memory
16 that is addressable; and

17 display transfer facility that writes the data generated from the
18 network to the increased memory allocated to the network generated data and
19 transfers the increased memory contents to the video display system so that the
20 network generated data is displayed along with the operating system user
21 interface.

1 47. The system of claim 46 wherein an online connection to
2 the network is reestablished automatically when the displayed network
3 generated data is accessed in a manner that indicates that the data is to be
4 updated from the network.

1 48. The system of claim 46 wherein the displayed network
2 generated data is a web page from the Internet.

1 49. The system of claim 46 wherein the displayed network
2 generated data includes an input field associated with a search engine, the
3 system further comprising:
4 network search facility that
5 receives a network search request that has been entered
6 into the input field;
7 sends the received search request to the search engine;
8 receives search results from the search engine; and
9 displays the search results on the display area allocated for
10 the network generated data.

1 50. The system of claim 46 wherein the displayed network
2 generated data includes an input field, the system further comprising:
3 network browse facility that
4 receives an address of additional network data that has
5 been entered into the input field;
6 sends the received address over the network to retrieve the
7 additional network data;
8 receives the additional data; and
9 displays the received data on the display area allocated for
10 the network generated data.

1 51. The system of claim 46 wherein the display transfer
2 facility writes at least a portion of the network generated data to the additional
3 memory such that the portion of data is displayed along with the operating
4 system user interface in a manner that prohibits the operating system user
5 interface from overwriting the portion of data.

1 52. The system of claim 46 wherein the display resolution
2 adjuster increases the displayable area of the video display system to include a
3 new display area portion by changing the resolution of the video display system
4 and by reducing the resolution available to the operating system user interface.





FIG. 3




FIG. 4




FIG. 5





FIG. 6




FIG. 7

FIG. 8

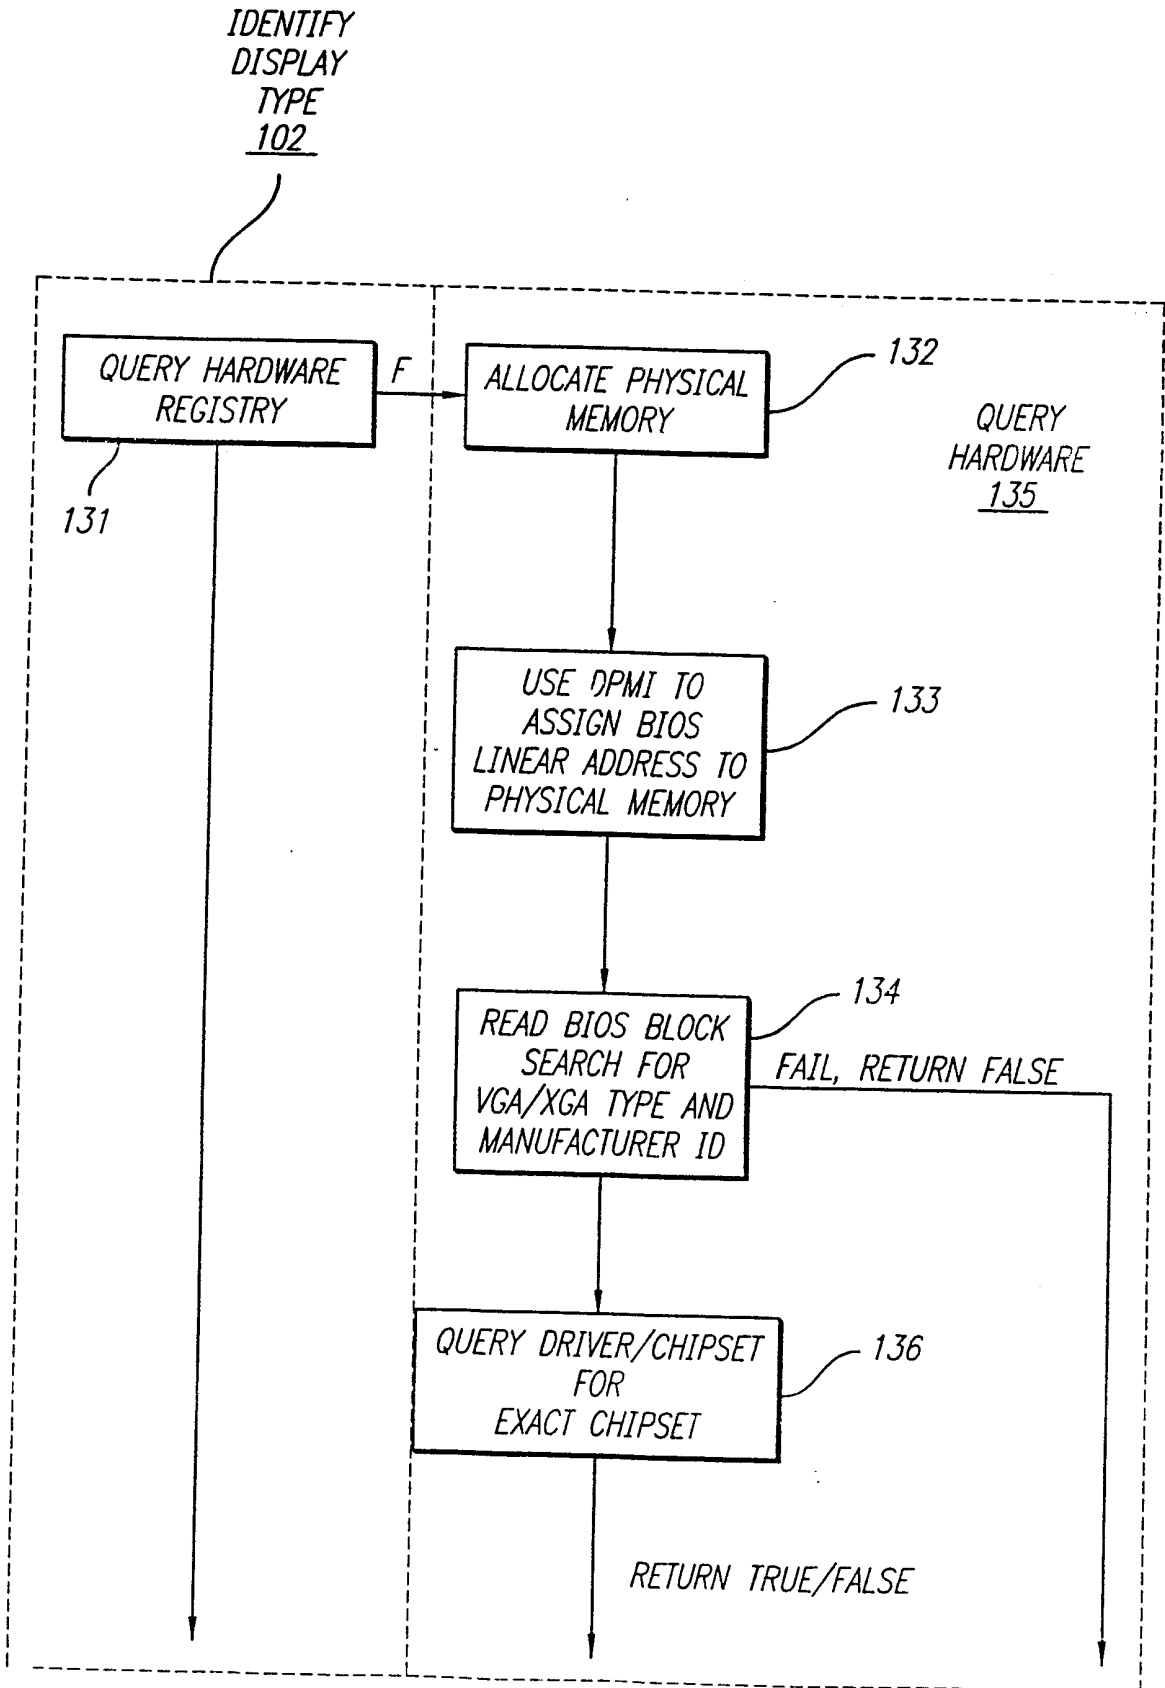
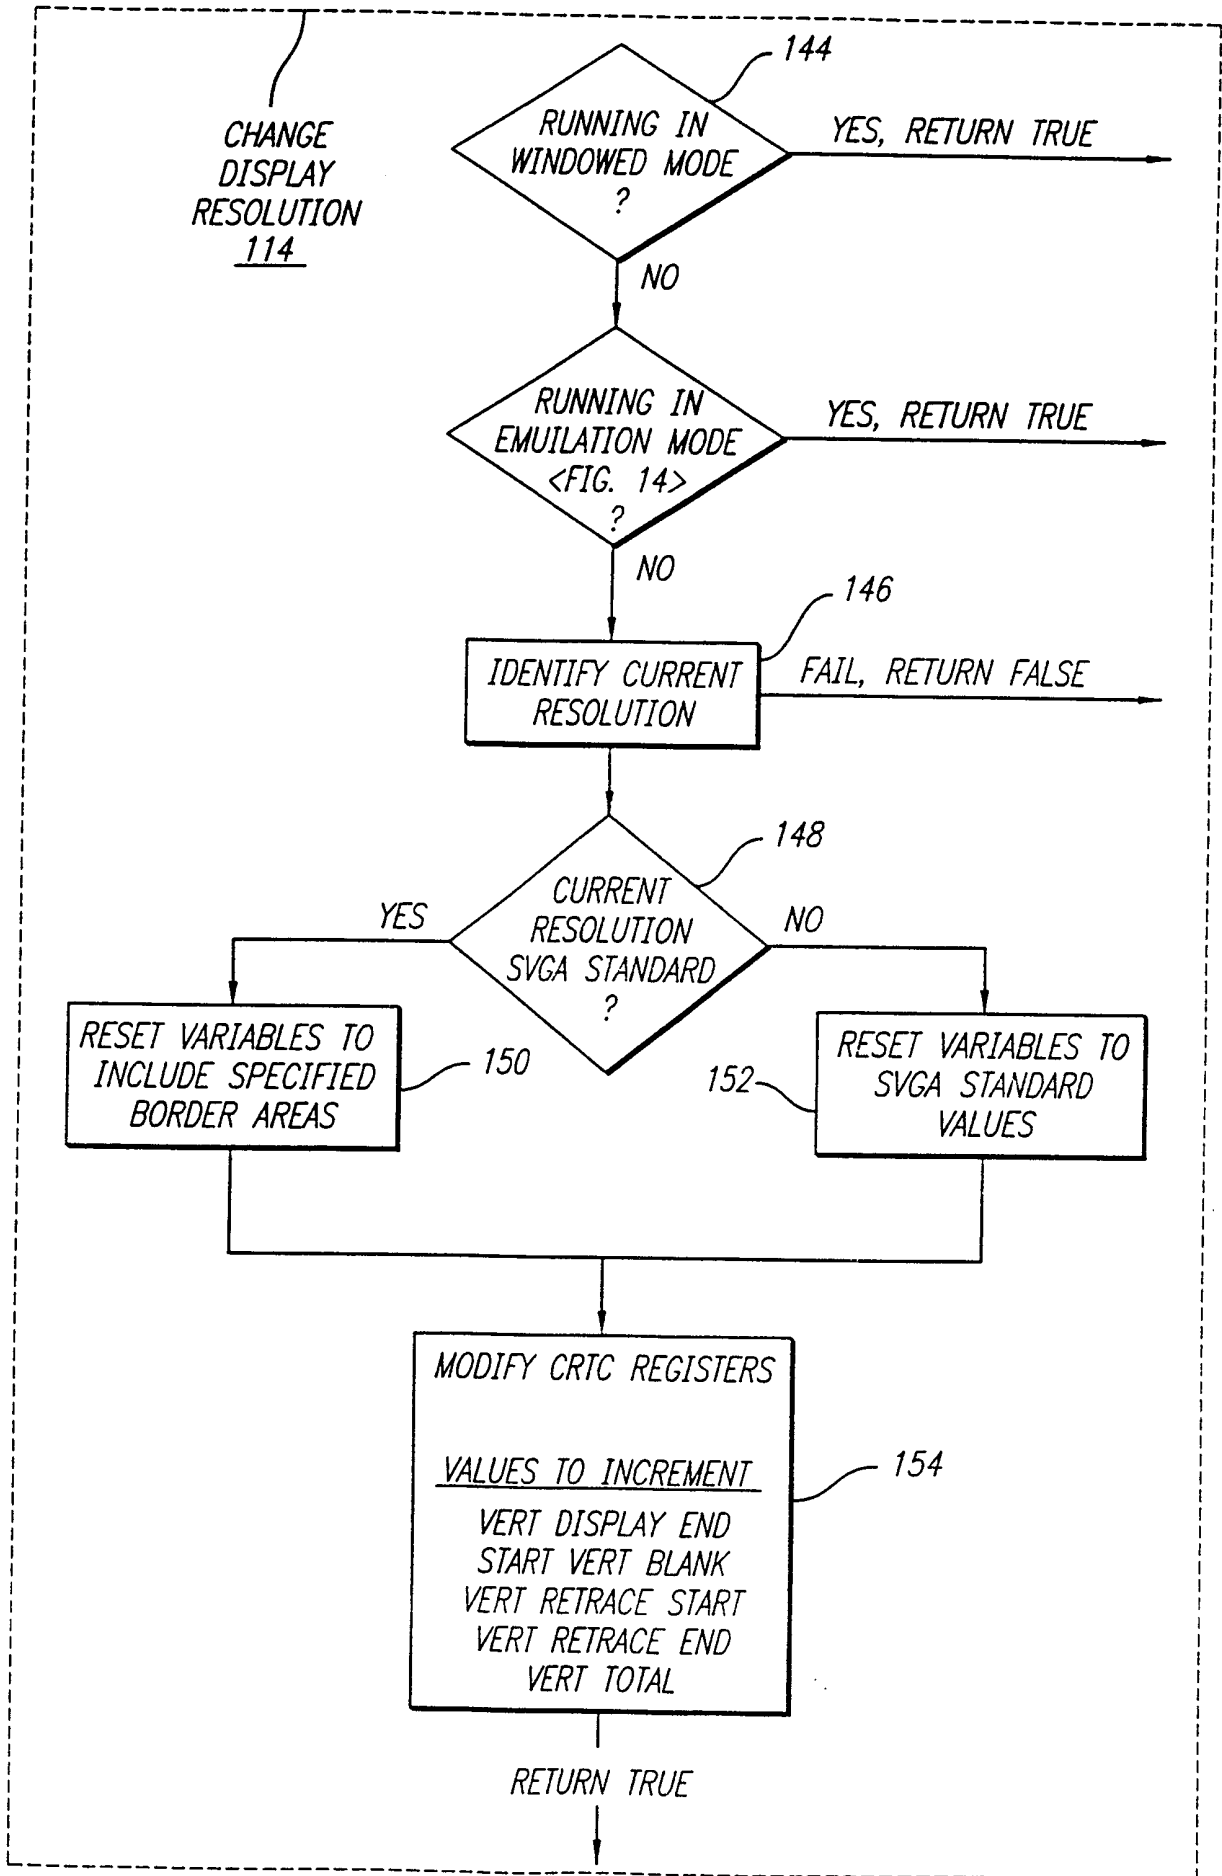


FIG. 9






FIG. 11

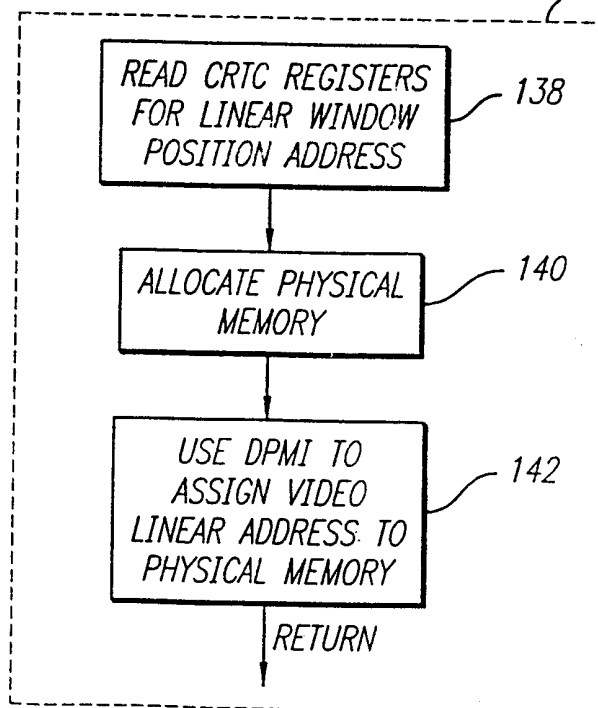


FIG. 12

MESSAGE PROCESS LOOP
USER INTERFACE

122

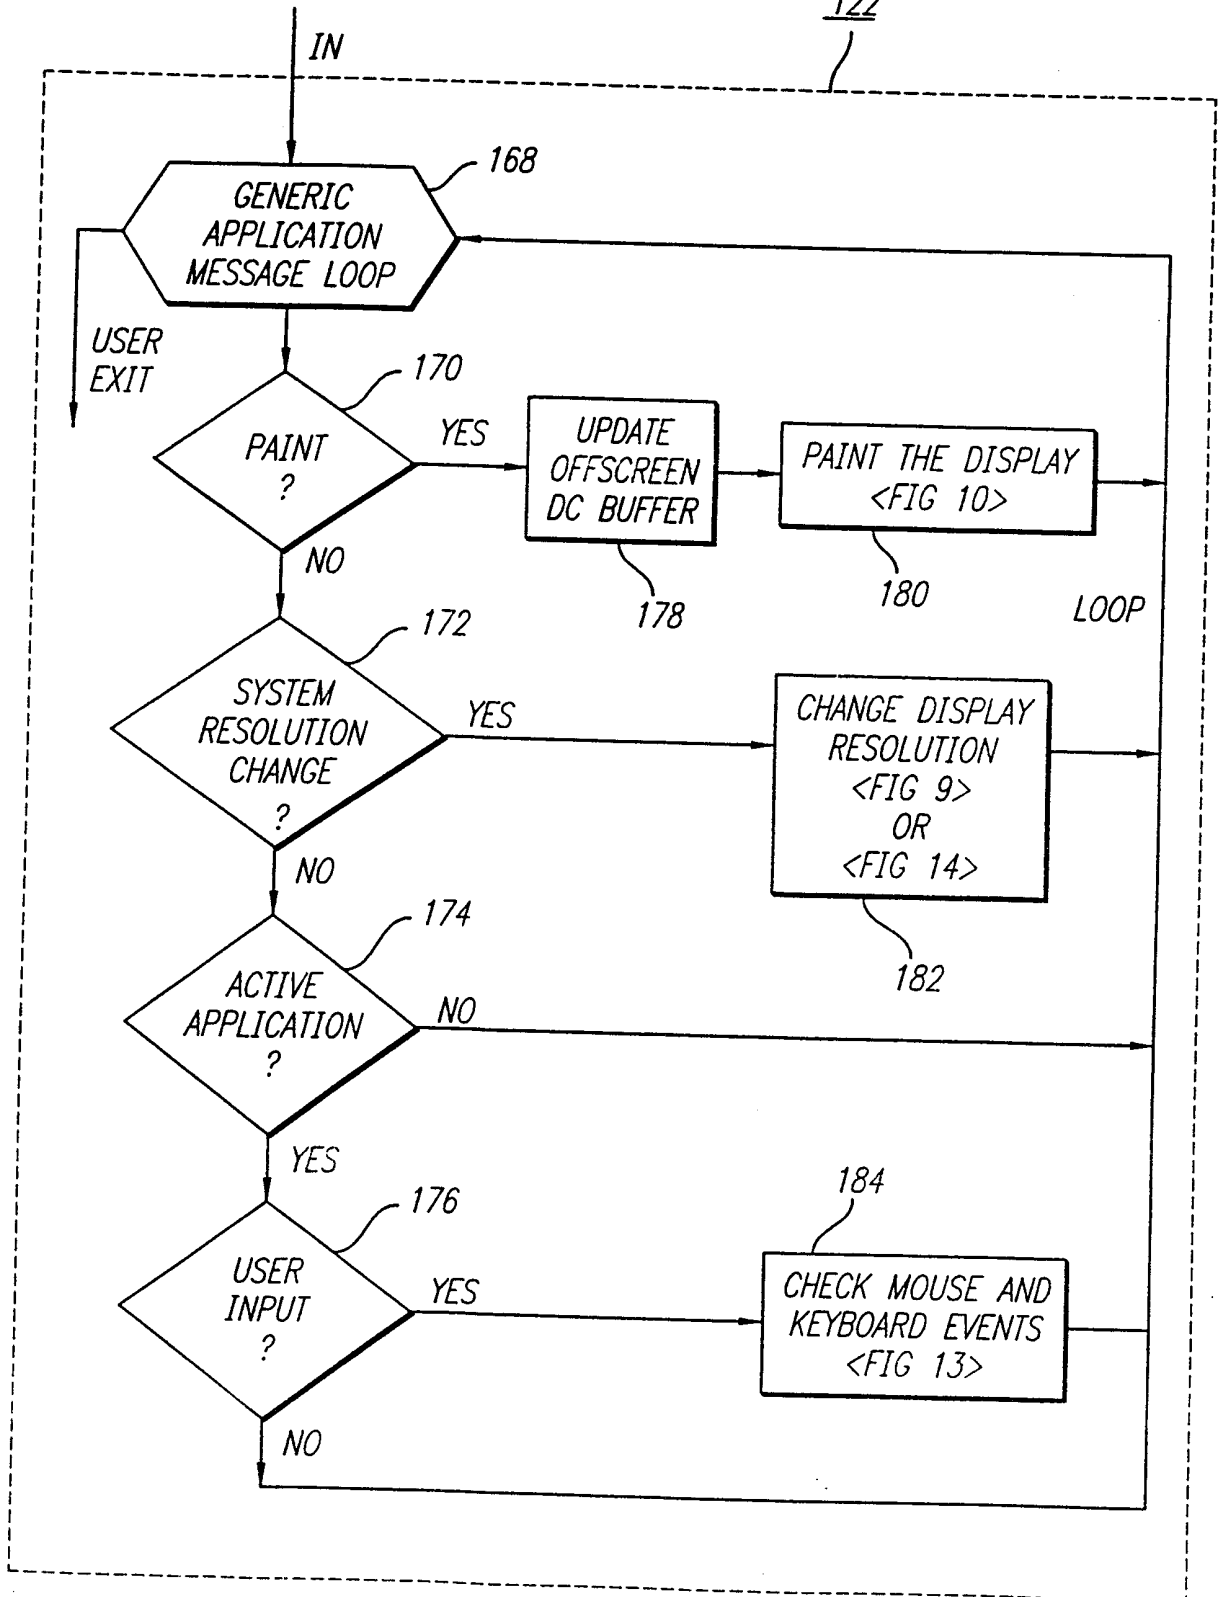


FIG. 13

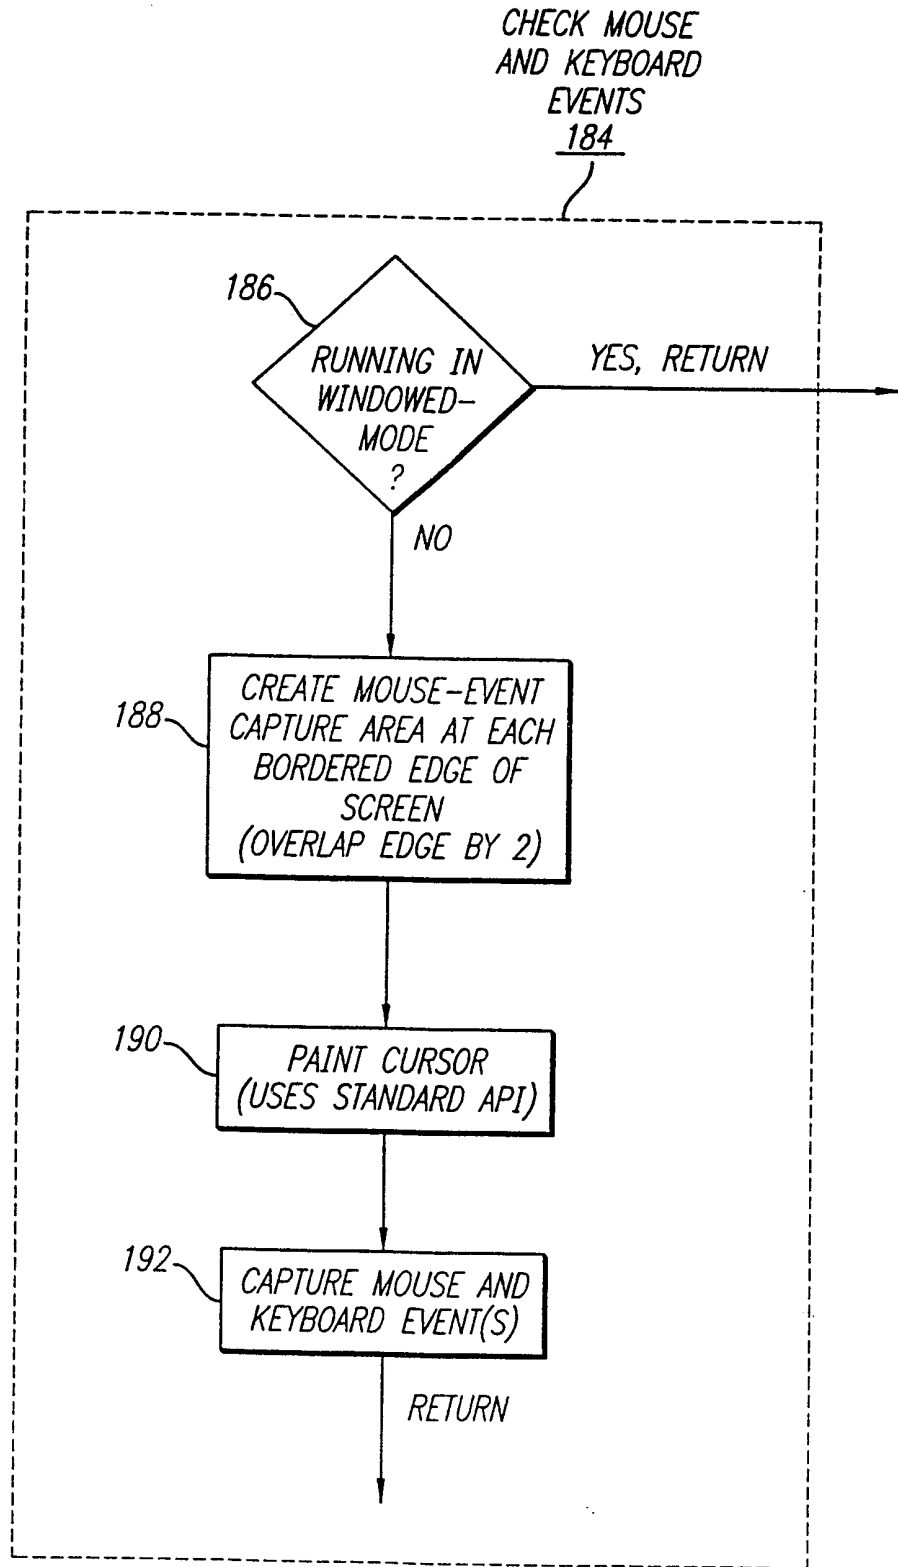


FIG. 14

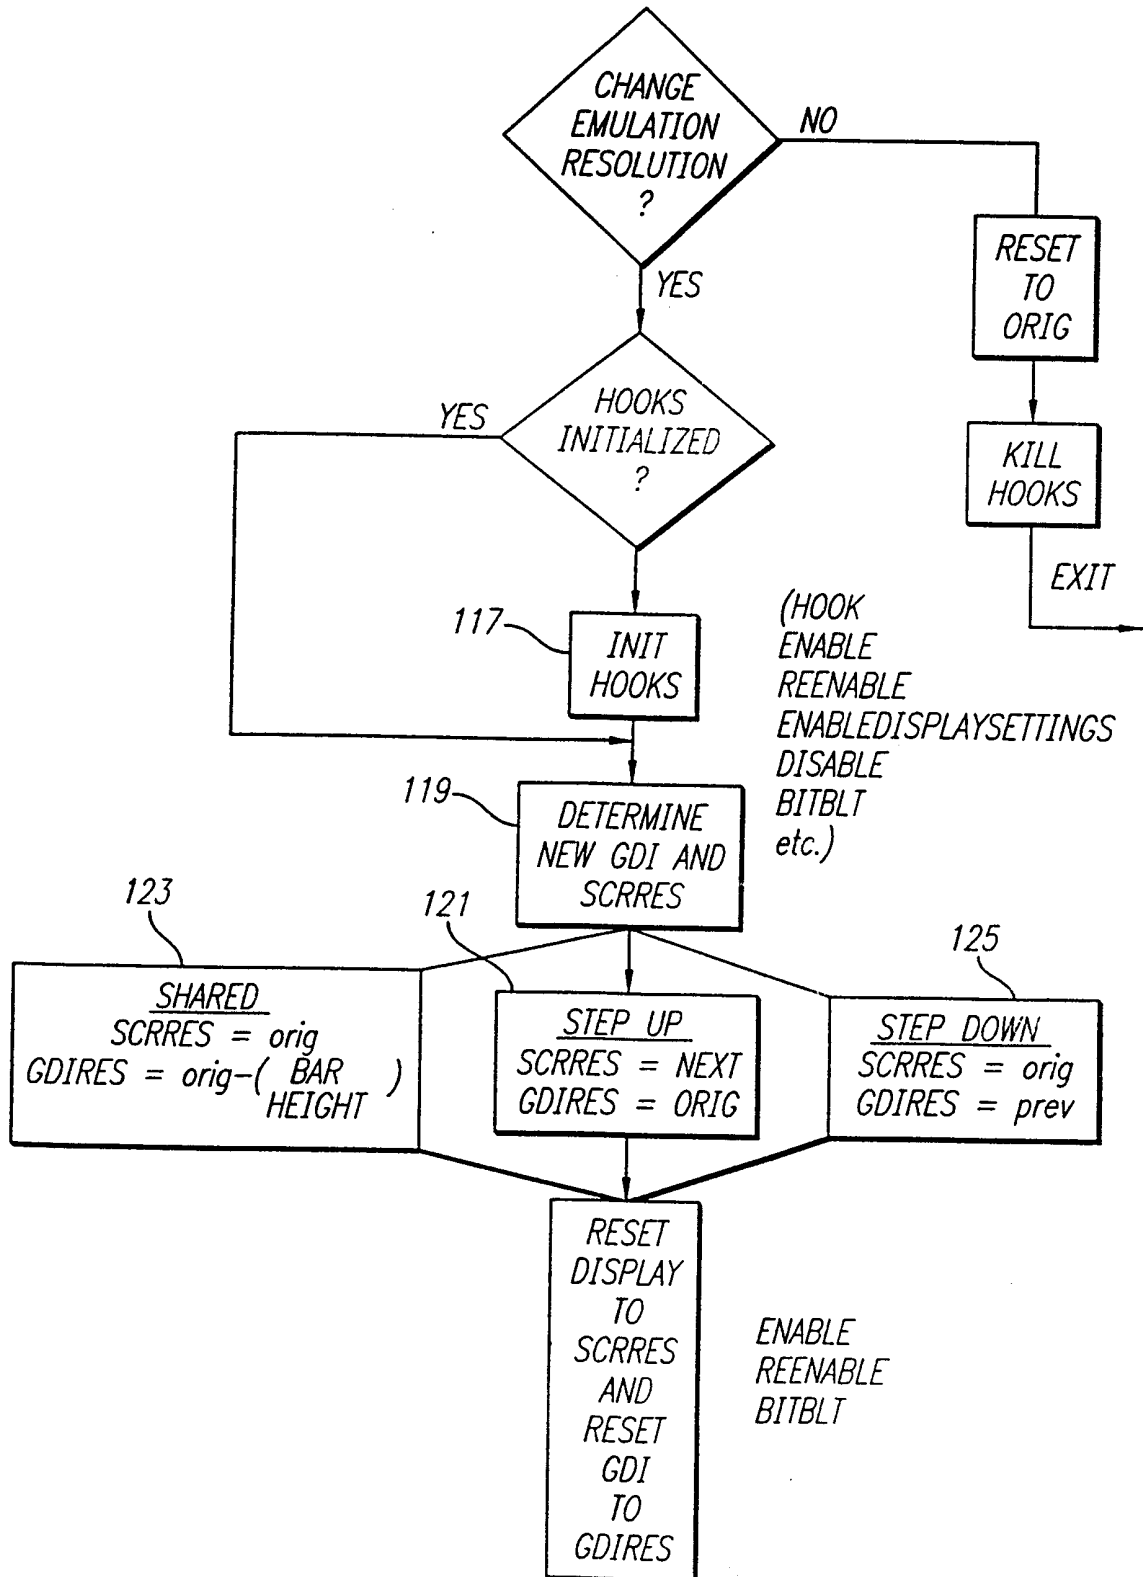


FIG. 15
PRIOR ART





FIG. 16



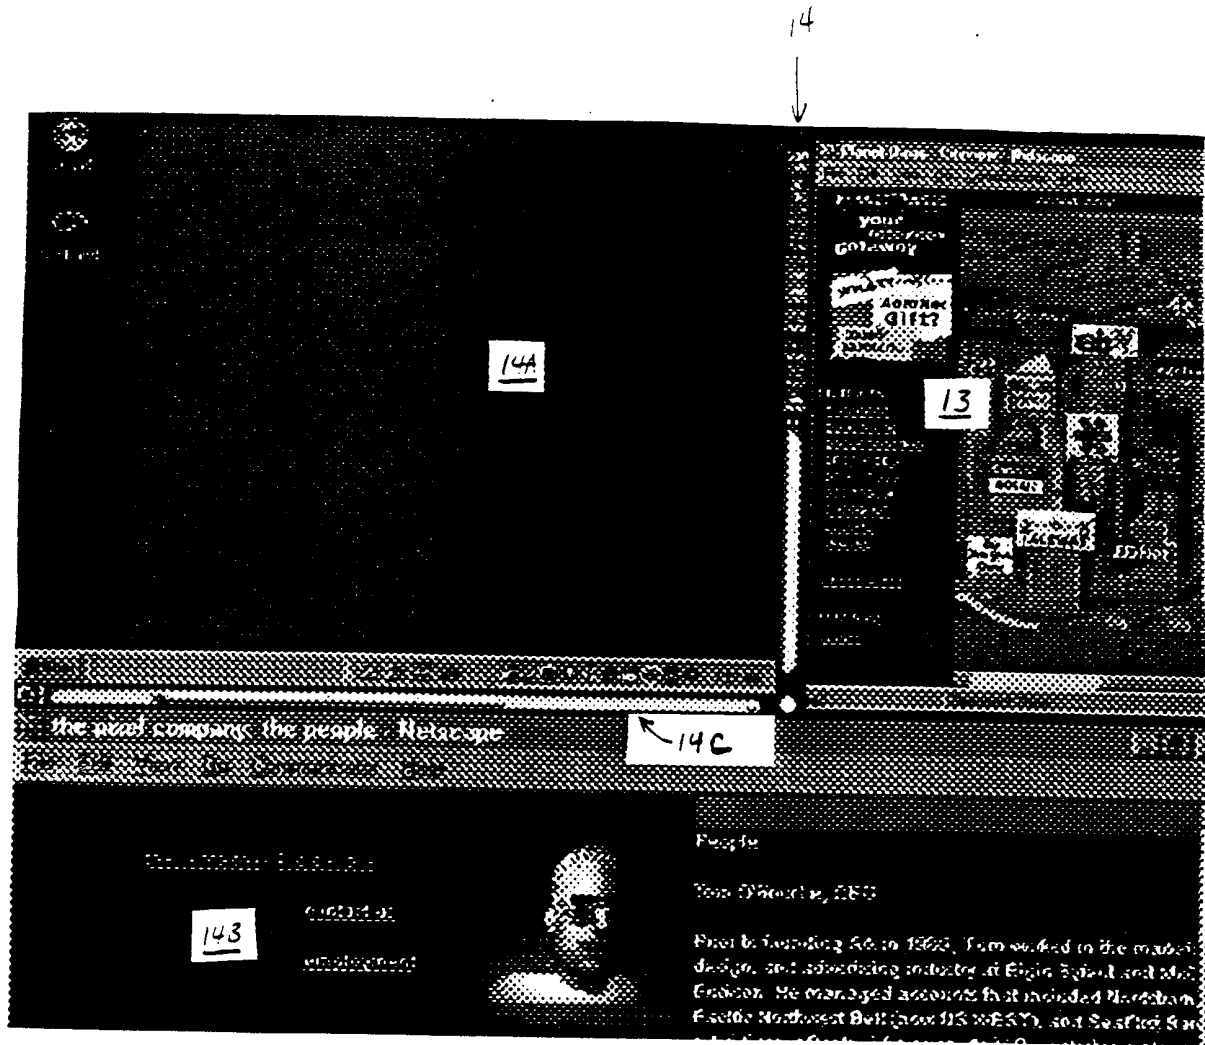


Fig 17




Fig. 18