




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification ⁶ : G06F 9/44</p>	<p>A1</p>	<p>(11) International Publication Number: WO 95/18410 (43) International Publication Date: 6 July 1995 (06.07.95)</p>
<p>(21) International Application Number: PCT/US94/10195 (22) International Filing Date: 12 September 1994 (12.09.94) (30) Priority Data: 08/175,851 30 December 1993 (30.12.93) US (71) Applicant: TALIGENT, INC. [US/US]; 10201 N. de Anza Boulevard, Cupertino, CA 95014 (US). (72) Inventors: ORTON, Debra, L.; 972 Nantucket Court, San Jose, CA 95126 (US). ROLLIN, Keith, A.; 1599 Meadowlark Lane, Sunnyvale, CA 94087 (US). GIBBONS, Bill; 24737 Prospect Avenue, Los Altos Hills, CA 95022 (US). (74) Agent: STEPHENS, L., Keith; Taligent, Inc., 10201 N. De Anza Boulevard, Cupertino, CA 95014 (US).</p>	<p>(81) Designated States: AT, AU, BB, BG, BR, BY, CA, CH, CN, CZ, DE, DK, ES, FI, GB, HU, JP, KP, KR, KZ, LK, LU, LV, MG, MN, MW, NL, NO, NZ, PL, PT, RO, RU, SD, SE, SK, UA, UZ, VN, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG).</p> <p>Published <i>With international search report.</i></p>	

(54) Title: OBJECT-ORIENTED INVOCATION FRAMEWORK

(57) Abstract

A view system is disclosed which provides support polymorphic initialization and finalization of objects. This mechanism is used by the view system to support virtual method invocation at construction/destruction time for objects that require the needed behavior. The C++ language does not support the invocation of virtual methods from within the constructor/destructor of a given object. The view system provides a mechanism whereby an initialization method is invoked after the object construction is completed, but before any methods are invoked on the constructed object. A similar mechanism is used to invoke a virtual finalization method before the object destruction is started. This mechanism is not view system-specific, but is provided and used by the view system to support such behavior.



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.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgystan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

OBJECT-ORIENTED INVOCATION FRAMEWORK

COPYRIGHT NOTIFICATION

Portions of this patent application contain materials that are subject to
5 copyright protection. The copyright owner has no objection to the facsimile
reproduction by anyone of the patent document, or the patent disclosure, as it
appears in the Patent and Trademark Office.

CROSS REFERENCE TO RELATED PATENT APPLICATIONS

10 This patent application is related to the patent application entitled Object
Oriented Area System, by Richard Daniel Webb et al. filed 6/20/93, and assigned
to Taligent, the disclosure of which is hereby incorporated by reference.

15 Field of the Invention

This invention generally relates to improvements in computer systems
and, more particularly, to operating system software for managing drawing
areas, called views, inside of a window display area in a graphic user interface.

20 Background of the Invention

One of the most important aspects of a modern computing system is the
interface between the human user and the machine. The earliest and most
popular type of interface was text based; a user communicated with the machine
by typing text characters on a keyboard and the machine communicated with
25 the user by displaying text characters on a display screen. More recently, graphic
user interfaces have become popular where the machine communicates with a
user by displaying graphics, including text and pictures, on a display screen and
the user communicates with the machine both by typing in textual commands
and by manipulating the displayed pictures with a pointing device, such as a
30 mouse.

Many modern computer systems operate with a graphic user interface
called a window environment. In a typical window environment, the graphical
display portrayed on the display screen is arranged to resemble the surface of an
electronic "desktop" and each application program running on the computer is
35 represented as one or more electronic "paper sheets" displayed in rectangular
regions of the screen called "windows".

Each window region generally displays information which is generated
by the associated application program and there may be several window regions

simultaneously present on the desktop, each representing information generated by a different application program. An application program presents information to the user through each window by drawing or "painting" images, graphics or text within the window region. The user, in turn, communicates with the application by "pointing at" objects in the window region with a cursor which is controlled by a pointing device and manipulating or moving the objects and also by typing information into the keyboard. The window regions may also be moved around on the display screen and changed in size and appearance so that the user can arrange the desktop in a convenient manner.

Each of the window regions also typically includes a number of standard graphical objects such as sizing boxes, buttons and scroll bars. These features represent user interface devices that the user can point at with the cursor to select and manipulate. When the devices are selected or manipulated, the underlying application program is informed, via the window system, that the control has been manipulated by the user.

In general, the window environment described above is part of the computer operating system. The operating system also typically includes a collection of utility programs that enable the computer system to perform basic operations, such as storing and retrieving information on a disc memory and performing file operations including the creation, naming and renaming of files and, in some cases, performing diagnostic operations in order to discover or recover from malfunctions.

The last part of the computing system is the "application program" which interacts with the operating system to provide much higher level functionality, perform a specific task and provide a direct interface with the user. The application program typically makes use of operating system functions by sending out series of task commands to the operating system which then performs a requested task, for example, the application program may request that the operating system store particular information on the computer disc memory or display information on the video display.

Figure 1 is a schematic illustration of a typical prior art computer system utilizing both an application program and an operating system. The computer system is schematically represented by box 100, the application is represented by box 102 and the operating system by box 106. The previously-described interaction between the application program 102 and the operating system 106 is illustrated schematically by arrow 104. This dual program system is used on many types of computer systems ranging from main frames to personal computers.

The method for handling drawing to screen displays varies from computer to computer and, in this regard, Figure 1 represents a prior art personal computer system. In order to provide drawing to screen displays, application program 102 generally stores information to be displayed (the storing operation is shown schematically by arrow 108) into a screen buffer 110. Under control of various hardware and software in the system the contents of the screen buffer 110 are read out of the buffer and provided, as indicated schematically by arrow 114, to a display adapter 112. The display adapter 112 contains hardware and software (sometimes in the form of firmware) which converts the information in screen buffer 110 to a form which can be used to drive the display monitor 118 which is connected to display adapter 112 by cable 116.

The prior art configuration shown in Figure 1 generally works well in a system where a single application program with a single thread of execution 102 is running at any given time. This simple system works properly because the single application program thread 102 can write information into any area of the entire screen buffer area 110 without causing a display problem. However, if the configuration shown in Figure 1 is used in a computer system where more than one application program, or more than one thread of execution in that application program 102 can be operational at the same time (for example, a "multitasking" computer system) display problems can arise. More particularly, if each thread in each application program has access to the entire screen buffer 110, in the absence of some direct communication between threads and applications, one thread may overwrite a portion of the screen buffer which is being used by another thread, thereby causing the display generated by one thread to be overwritten by the display generated by the other thread.

Accordingly, mechanisms were developed to coordinate the operation of the applications as well as the threads of execution within the applications to ensure that each application thread was confined to only a portion of the screen buffer thereby separating the other displays. This coordination became complicated in systems where multiple "windows" with multiple threads drawing to those windows were allowed. The problem was divided into two pieces: Managing the windows and their display area (application programs) and managing the threads of execution within those applications. The Window Manager handles coordination between applications and their windows. The View System handles coordination of threads within the applications and their window(s). Each window is subdivided in a hierarchy of drawing areas called

"views" which are associated with specific threads of execution within a given application program.

When the application window is arranged so that views appear to "overlap", a view which appears in the window in "front" of another view covers and obscures part of the underlying view. Thus, except for the foremost view, only part of the underlying views may be drawn on the screen and be "visible" at any given time. Further, because the view can be moved or resized by the user, the portion of each view which is visible changes as other views are added, removed, moved or resized. Thus, the portion of the window which is assigned to each thread also changes as views from other threads are added, removed, moved or resized.

In order to efficiently manage the changes to the window necessary to accommodate rapid screen changes caused by moving or resizing views, the prior art computer arrangement shown in Figure 1 was modified as shown in Figure 2. In this new arrangement computer system 200 is controlled by one or more application programs, comprised of one or more threads of execution, of which threads 202 and 216 are shown, and which may be running simultaneously in the computer system. Each of the threads interfaces with the operating system 204 as illustrated schematically by arrows 206 and 220. However, in order to display information on the display screen, application threads 202 and 216 send display information to a central View System 218 located in the application program 204. The view system 218, in turn, interfaces directly with the screen buffer 210 as illustrated schematically by arrow 208. The contents of screen buffer 210 are provided, as indicated by arrow 212, to a display adapter 214 which is connected by a cable 222 to a display monitor 224.

In such a system, the view system 218 is generally responsible for maintaining all of the display contents that the user sees within a window during operation of the application programs. Since the view system 218 is in communication with all the threads within an application, it can coordinate between threads to insure that view displays do not overlap. Consequently, it is generally the task of the view system to keep track of the location and size of the view and the view areas which must be drawn and redrawn as views and windows are moved.

The view system 218 receives display requests from each of the application threads 202 and 216. However, since only the view system 218 interfaces with the screen buffer 210, it can allocate respective areas of the screen buffer 210 for each application and insure that no thread erroneously overwrites the display generated by another thread. There are a number of different

window environments commercially available which utilize the arrangement illustrated in Figure 2. These include the X/Window Operating environment, the WINDOWS graphical user interface developed by the Microsoft Corporation and the OS/2 Presentation Manager developed by the International Business Machines Corporation.

Each of these window environments has its own internal software architecture, but the architectures can all be classified by using a multi-layer model similar to the multi-layer models used to describe computer network software. A typical multi-layer model includes the following layers:

User Interface
Window Manager
Resource Control and Communication
Component Driver Software
Computer Hardware

where the term "window environment" refers to all of the above layers taken together.

The lowest or computer hardware level includes the basic computer and associated input and output devices including display monitors, keyboards, pointing devices, such as mice or trackballs, and other standard components, including printers and disc drives. The next or "component driver software" level consists of device-dependent software that generates the commands and signals necessary to operate the various hardware components. The resource control and communication layer interfaces with the component drivers and includes software routines which allocate resources, communicate between applications and multiplex communications generated by the higher layers to the underlying layers. The view system handles the user interface to basic drawing operations, such as moving and resizing views, activating or inactivating views and redrawing and repainting views. The final user interface layer provides high level facilities that implement the various controls (buttons, sliders, boxes and other controls) that application programs use to develop a complete user interface.

Although the arrangement shown in Figure 2 solves the display screen interference problem, it suffers from the drawback that the view system must process the screen display requests generated by all of the application threads. Since the requests can only be processed serially, the requests are queued for presentation to the view system before each request is processed to generate a display on terminal 224. In a display where many views are present simultaneously on the screen, the view system 218 can easily become a

"bottleneck" for display information and prevent rapid changes of the display by the application threads 202 and 216. A delay in the redrawing of the screen when views are moved or repositioned by the user often manifests itself by the appearance that the views and windows are being constructed in a piecemeal fashion which becomes annoying and detracts from the operation of the system.

Accordingly, it is an object of the present invention to provide a view system which can interface with application threads in such a manner that the screen display generated by each application thread can be quickly and effectively redrawn.

It is another object of the present invention to provide a view system which coordinates the display generation for all of the application threads in order to prevent the applications from interfering with each other or overwriting each other on the screen display.

It is yet another object of the present invention to provide a view system which can interact with the application threads by means of a simple command structure without the application threads being concerned with actual implementation details.

It is yet another object of the present invention to provide a view system which allows application developers who need detailed control over the screen display process to achieve this control by means of a full set of display control commands which are available, but need not be used by each application thread.

It is yet another object of the present invention to provide a view system which provides application developers with a powerful and flexible drawing environment which includes a virtual coordinate space, arbitrarily shaped views (and windows) and up-to-date drawing state information to facilitate rapid, accurate drawing from multiple threads of execution.

It is yet another object of the present invention to provide a view system which provides application developers with an automatic system for keeping the display buffer up-to-date.

Summary of the Invention

The foregoing problems are overcome and the foregoing objects are achieved in an illustrative embodiment of the invention in which an object-oriented viewing framework provides support polymorphic initialization and finalization of objects. This mechanism is used by the view system to support virtual method invocation at construction/destruction time for objects that require the needed behavior. The C++ language does not support the invocation of virtual methods from within the constructor/destructor of a

given object. The view system provides a mechanism whereby an initialization method is invoked after the object construction is completed, but before any methods are invoked on the constructed object. A similar mechanism is used to invoke a virtual finalization method before the object destruction is started.

5 This mechanism is not view system-specific, but is provided and used by the view system to support such behavior.

Brief Description of the Drawings

The above and further advantages of the invention may be better understood by referring to the following description in conjunction with the accompanying drawings, in which:

Figure 1 is a schematic block diagram of a prior art computer system showing the relationship of the application thread, the operating system, the screen buffer and, the display monitor.

15 Figure 2 is a schematic block diagram of a modification of the prior art system shown in Figure 1 which allows several application thread threads running simultaneously to generate display output in a single window.

Figure 3 is a block schematic diagram of a computer system for example, a personal computer system on which the inventive object-oriented viewing framework operates.

Figure 4 is a schematic block diagram of a modified computer system showing the interaction between a plurality of application threads, the viewing framework, the window manager, and the screen buffer in order to display graphic information on the display monitor.

25 Figure 5 is a block schematic diagram of the information paths which indicate the manner in which an application thread communicates with the inventive object-oriented viewing framework and then directly to the screen buffer.

Figure 6 is a schematic diagram indicating the typical appearance of a graphical user interface which is supported by a object-oriented viewing framework illustrating the components and parts of a window and the view hierarchy which is contained within.

Figures 7A and 7B illustrate the portions of the view hierarchy which must be redrawn when a view is resized.

35 Figure 8 is an illustrative flow chart of a method by which the object-oriented viewing framework supports updating of invalid areas on the display by means of a background thread(s) in order to display information on the display screen.

Figure 9 is an illustrative flow chart of a method used by the application thread to create a new view object and install it in the view hierarchy.

Figure 10 is an illustrative flow chart of a method used to support polymorphic initialization and finalization of a view framework object.

5 Figure 11 is an illustrative flow chart of the method by which an application thread requests drawing state information from the object-oriented viewing framework.

10 Figure 12 is a block schematic diagram of two non-rectilinear views and a disjoint view, within the view hierarchy, which are supported by the object-oriented viewing framework.

Figure 13 is a block schematic diagram of the multiple coordinate spaces support by the object-oriented viewing framework.

Figure 14 is a block schematic diagram of aligned view objects based on their center of gravity specification by the object-oriented viewing framework.

15 Figure 15 is an illustrative flow chart of the method by which a view is automatically, spatially laid-out by the object-oriented viewing framework.

Figure 16 is a block schematic diagram of the visual effect, "Magnifier View", that uses a matrix object to apply a scaling transformation to the view below it in the view hierarchy as supported by the object-oriented viewing
20 framework.

Figure 17 is a block schematic diagram of the method used by the object-oriented viewing framework to enable grouping of related windows to be moved as a single layer.

25 Figure 18 is an illustrative flow chart of the method by which an application thread makes use of a non-multitasking aware object in a multitasking environment as provided by the object-oriented viewing framework

30 Figure 19 is an illustrative flow chart of the method by which an application thread receives a positional event via the input system and the object-oriented viewing framework.

Figure 20 is an illustrative flow chart of the method by which an application thread receives changes in the object-oriented viewing framework in a single batch notification.

35 Figure 21 is a block schematic diagram of the method used by the object-oriented viewing framework to support read-only and read-write operations on the hierarchy and view objects.

Detailed Description of a Preferred Embodiment of the Invention

The invention is preferably practiced in the context of an operating system resident on a personal computer such as the IBM PS/2 or Apple Macintosh computer. A representative hardware environment is depicted in Figure 3, which illustrates a typical hardware configuration of a computer 300 in accordance with the subject invention. The computer 300 is controlled by a central processing unit 302 (which may be a conventional microprocessor) and a number of other units, all interconnected via a system bus 308, are provided to accomplish specific tasks. Although a particular computer may only have some of the units illustrated in Figure 3, or may have additional components not shown, most computers will include at least the units shown.

Specifically, computer 300 shown in Figure 3 includes a random access memory (RAM) 306 for temporary storage of information, a read only memory (ROM) 304 for permanent storage of the computer's configuration and basic operating commands and an input/output (I/O) adapter 310 for connecting peripheral devices such as a disk unit 313 and printer 314 to the bus 308, via cables 315 and 312, respectively. A user interface adapter 316 is also provided for connecting input devices, such as a keyboard 320, and other known interface devices including mice, speakers and microphones to the bus 308. Visual output is provided by a display adapter 318 which connects the bus 308 to a display device 322, such as a video monitor. The workstation has resident thereon and is controlled and coordinated by operating system software such as the Apple System/7 operating system.

In a preferred embodiment, the invention is implemented in the C++ programming language using object-oriented programming techniques. C++ is a compiled language, that is, programs are written in a human-readable script and this script is then provided to another program called a compiler which generates a machine-readable numeric code that can be loaded into, and directly executed by, a computer. As described below, the C++ language has certain characteristics which allow a software developer to easily use programs written by others while still providing a great deal of control over the reuse of programs to prevent their destruction or improper use. The C++ language is well-known and many articles and texts are available which describe the language in detail. In addition, C++ compilers are commercially available from several vendors including Borland International, Inc. and Microsoft Corporation. Accordingly, for reasons of clarity, the details of the C++ language and the operation of the C++ compiler will not be discussed further in detail herein.

As will be understood by those skilled in the art, Object-Oriented Programming (OOP) techniques involve the definition, creation, use and

destruction of "objects". These objects are software entities comprising data elements and routines, or functions, which manipulate the data elements. The data and related functions are treated by the software as an entity and can be created, used and deleted as if they were a single item. Together, the data and functions enable objects to model virtually any real-world entity in terms of its characteristics, which can be represented by the data elements, and its behavior, which can be represented by its data manipulation functions. In this way, objects can model concrete things like people and computers, and they can also model abstract concepts like numbers or geometrical designs.

Objects are defined by creating "classes" which are not objects themselves, but which act as templates that instruct the compiler how to construct the actual object. A class may, for example, specify the number and type of data variables and the steps involved in the functions which manipulate the data. An object is actually created in the program by means of a special function called a constructor which uses the corresponding class definition and additional information, such as arguments provided during object creation, to construct the object. Likewise objects are destroyed by a special function called a destructor. Objects may be used by using their data and invoking their functions.

The principle benefits of object-oriented programming techniques arise out of three basic principles; encapsulation, polymorphism and inheritance. More specifically, objects can be designed to hide, or encapsulate, all, or a portion of, the internal data structure and the internal functions. More particularly, during program design, a program developer can define objects in which all or some of the data variables and all or some of the related functions are considered "private" or for use only by the object itself. Other data or functions can be declared "public" or available for use by other programs. Access to the private variables by other programs can be controlled by defining public functions for an object which access the object's private data. The public functions form a controlled and consistent interface between the private data and the "outside" world. Any attempt to write program code which directly accesses the private variables causes the compiler to generate an error during program compilation which error stops the compilation process and prevents the program from being run.

Polymorphism is a concept which allows objects and functions which have the same overall format, but which work with different data, to function differently in order to produce consistent results. For example, an addition function may be defined as variable A plus variable B ($A+B$) and this same

format can be used whether the A and B are numbers, characters or dollars and cents. However, the actual program code which performs the addition may differ widely depending on the type of variables that comprise A and B.

Polymorphism allows three separate function definitions to be written, one for each type of variable (numbers, characters and dollars). After the functions have been defined, a program can later refer to the addition function by its common format (A+B) and, during compilation, the C++ compiler will determine which of the three functions is actually being used by examining the variable types. The compiler will then substitute the proper function code.

Polymorphism allows similar functions which produce analogous results to be "grouped" in the program source code to produce a more logical and clear program flow.

The third principle which underlies object-oriented programming is inheritance, which allows program developers to easily reuse pre-existing programs and to avoid creating software from scratch. The principle of inheritance allows a software developer to declare classes (and the objects which are later created from them) as related. Specifically, classes may be designated as subclasses of other base classes. A subclass "inherits" and has access to all of the public functions of its base classes just as if these function appeared in the subclass. Alternatively, a subclass can override some or all of its inherited functions or may modify some or all of its inherited functions merely by defining a new function with the same form (overriding or modification does not alter the function in the base class, but merely modifies the use of the function in the subclass). The creation of a new subclass which has some of the functionality (with selective modification) of another class allows software developers to easily customize existing code to meet their particular needs.

Although object-oriented programming offers significant improvements over other programming concepts, program development still requires significant outlays of time and effort, especially if no pre-existing software programs are available for modification. Consequently, a prior art approach has been to provide a program developer with a set of pre-defined, interconnected classes which create a set of objects and additional miscellaneous routines that are all directed to performing commonly-encountered tasks in a particular environment. Such pre-defined classes and libraries are typically called "frameworks" and essentially provide a pre-fabricated structure for a working application.

For example, a framework for a user interface might provide a set of pre-defined graphic interface objects which create windows, scroll bars, menus, etc.

and provide the support and "default" behavior for these graphic interface objects. Since frameworks are based on object-oriented techniques, the pre-defined classes can be used as base classes and the built-in default behavior can be inherited by developer-defined subclasses and either modified or overridden to allow developers to extend the framework and create customized solutions in a particular area of expertise. This object-oriented approach provides a major advantage over traditional programming since the programmer is not changing the original program, but rather extending the capabilities of the original program. In addition, developers are not blindly working through layers of code because the framework provides architectural guidance and modeling and, at the same time, frees the developers to supply specific actions unique to the problem domain.

There are many kinds of frameworks available, depending on the level of the system involved and the kind of problem to be solved. The types of frameworks range from high-level application frameworks that assist in developing a user interface, to lower-level frameworks that provide basic system software services such as communications, printing, file systems support, graphics, etc. Commercial examples of application frameworks include MacApp (Apple), Bedrock (Symantec), OWL (Borland), NeXT Step App Kit (NeXT), and Smalltalk-80 MVC (ParcPlace).

While the framework approach utilizes all the principles of encapsulation, polymorphism, and inheritance in the object layer, and is a substantial improvement over other programming techniques, there are difficulties which arise. Application frameworks generally consist of one or more object "layers" on top of a monolithic operating system and even with the flexibility of the object layer, it is still often necessary to directly interact with the underlying operating system by means of awkward procedural calls.

In the same way that an application framework provides the developer with prefab functionality for an application thread, a system framework, such as that included in a preferred embodiment, can provide a prefab functionality for system level services which developers can modify or override to create customized solutions, thereby avoiding the awkward procedural calls necessary with the prior art application frameworks programs. For example, consider a display framework which could provide the foundation for creating, deleting and manipulating windows to display information generated by an application thread. An application software developer who needed these capabilities would ordinarily have to write specific routines to provide them. To do this with a framework, the developer only needs to supply the characteristics and behavior

of the finished display, while the framework provides the actual routines which perform the tasks.

A preferred embodiment takes the concept of frameworks and applies it throughout the entire system, including the application and the operating system. For the commercial or corporate developer, systems integrator, or OEM, this means all of the advantages that have been illustrated for a framework such as MacApp can be leveraged not only at the application level for such things as text and user interfaces, but also at the system level, for services such as printing, graphics, multi-media, file systems, I/O, testing, etc.

Figure 4 shows a schematic overview of a computer system utilizing the object-oriented viewing framework of the present invention. The computer system is shown generally as box 400 and an application 416 with multiple threads of execution (of which application threads 401 and 402 are shown) and an operating system 404 are provided to control and coordinate the operations of the computer. In order to simplify Figure 4, the interaction of the application 416 with the operating system 404 is limited to the interactions dealing with the screen displays. As shown in the figure, both application threads 401 and 402 interface with the view system portion 405 of the application program as illustrated by arrows 403 and 406. The view system 405, in turn, sends information to the screen buffer 410 and the operating system 404 as schematically illustrated by arrow 407 and 420.

However, in accordance with the invention, and, as shown in Figure 4, application threads 401 and 402 also directly send information to the screen buffer 410 as illustrated by arrow 407. As will hereinafter be explained in detail, application threads 401 and 402 provide display information directly to the screen buffer 410 and retrieve stored information from the view system 405 when drawing is required. More specifically, when a view is changed, the view system 405 recomputes and stores the visible area of each view. This stored visible area is retrieved by the respective application thread and used as a clipping region into which the application draws the display information. Repainting or drawing of the view is performed simultaneously by the application threads in order to increase the screen repainting speed.

The application displays are kept separated on the display screen because the view system 405 recomputes the view visible areas so that none of the areas overlap. Thus, if each application thread, such as application thread 401 or application thread 402 draws only in the visible area provided to it by the view system 405, there will be no overlap in the displays produced by the screen buffer. Once the display information is drawn into the screen buffer 410 it is

provided, as indicated by arrow 412, to a display adapter 414 which is, in turn, connected by cable, or bus, 422 to the display monitor 424.

The interaction of an application thread with the view system is illustrated in more detail in schematic diagram Figure 5. As previously mentioned, the view system (illustrated as box 510 in Figure 5) is an object-oriented program. Accordingly, an application thread 508 interfaces with the view system by creating and manipulating "objects". In particular, each application thread creates a view hierarchy object, for example, view hierarchy object 512 in order to communicate with view system 510. The application thread 508 then communicates with the view hierarchy object 512 by creating a view object 506 and installing it in the hierarchy as shown schematically by arrow 502. The view system itself is a collection of objects which is created when the application program is started. The view system 510 interfaces with the operating system 500 via a data stream 504 to perform window operations on behalf of the application program and view system 510.

As will hereinafter be described in more detail, each view object 506 includes a small data store or "cache" area, called the drawing state 514 which is used to store the associated view visible area and other drawing-related state (coordinate system etc.). When the application thread desires to redraw the information in one of its associated views, the view object first checks cache status. If the information stored in the cache has not been changed or invalidated, then this information is used to redraw the window. The use of the cache area reduces the time necessary to complete a redrawing operation.

Since many view objects may be created simultaneously in order to simultaneously display many views within a window, each view object 506 communicates with the view system 510 by means of multitask-safe method calls 502. The view system communicates with the operating system via data stream 504 by creating "stream" objects which contain the software commands necessary to transfer information from one object to another. For example, when operating system 500 desires to transfer information to view system object 510, operating system 500 creates a stream object which "streams" the data into view system object 510. Similarly, when view system object 510 desires to transfer information back to operating system 500, view system object 510 creates a stream object which "streams" the data into window object 500. Such stream objects are conventional in nature and not described in detail herein. The stream objects which carry data from operating system 500 to view system object 510 and the stream objects which carry information from view system object 510 to operating system 500 are illustrated collectively as arrow 504.

As shown in Figure 5, view system object 510 consists of four main parts: the view hierarchy framework 512, the one or more view objects (and their associated command objects) installed in the hierarchy 506, the drawing state cached for each view in the hierarchy 514, and the framework to support background updating of "invalid" areas within views installed in the hierarchy 516. The background drawing framework 516 includes an independent task which is started by the view system 510 when the view system 510 is created. As will be hereinafter explained in detail, the background updating framework is responsible for refreshing the portions of the views 506 in the view hierarchy 512 which are visible on the data display screen and have become invalid via view hierarchy changes, local view changes, or window manipulations. To this end, it collects areas in need of updating, recomputes a view's visible area intersected with the "invalid" area and tells the affected view to refresh only that portion of itself that is invalid.

The drawing state cache 514 is comprised of a number of different objects which together store information related to the drawing attributes and environment for the associated view. The drawing state cache is created by the view system on behalf of individual views and maintained by the view system including a "time stamp" indicating the time of the last modification.

As previously mentioned, in accordance with a preferred embodiment, the operating system is capable of running multiple threads simultaneously and, whenever two or more threads are operating simultaneously, there is a potential for mutual interaction. Such mutual interaction can occur when two or more threads attempt to access simultaneously shared resources, such as the view hierarchy. Accordingly, concurrency controls are necessary to manage such interactions and to prevent unwanted interference. An illustrative concurrency control technique known as a monitor is used in one embodiment. Monitors are well-known devices which are used to "serialize" concurrent access attempts to a resource and provide an object-oriented interface. In particular, before a thread can access a resource which is controlled by a monitor, the thread must "acquire" the monitor. When the thread is finished with the resource it releases the monitor for acquisition by another thread. Each monitor generally has a request queue associated with it so that requests to acquire the monitor which cannot be honored (because the monitor has been acquired by another thread) are held on the queue (called "blocking"). Monitors also provide a mechanism to prevent acquisition until a specified "condition" has been met.

In the present system, monitors are used to protect several different shared resources. In particular, a view system monitor is used to prevent the application threads from interacting with the view system simultaneously and corrupting data. Before accessing the cached drawing data, each application
5 thread must acquire the view system monitor. The monitor used in the present system allows individual threads to "reacquire" a given monitor to prevent a thread dead-locking with itself.

Figure 6 shows an illustrative screen display generated by a typical "window environment" program. A window 600 is an area enclosed by borders
10 which can be moved and resized in a conventional manner. The window 600 usually includes a title bar 606 and a menu bar 604, each of which are a view and may themselves contain other views. The menu bar 604, also a view allows access to a number of pull-down menus (not shown) that are operated in a well-known manner and allow the user to access various file, editing and other
15 commands.

The area remaining within the window, after excluding the title bar 606 , the menu bar 604 and the borders, is called the "content" view and constitutes the main area that can be drawn or painted by an application thread such as a drawing program. A content view may enclose additional views called "child"
20 views that are associated with the one or more application threads. In this case the containing view is called a "parent" view or "container" view in relation to the child view(s). Each child view may also have one or more child views associated with it for which it is a parent view and so on, thus constructing a view hierarchy.

Many application threads further sub-divide the content view into a
25 number of child views which are independently controlled. These typically include a document view 622, a "toolbar" or "palette" view 616, and, in some cases, a status line view (not shown). The document view 622 may be equipped with horizontal and vertical scroll bar views, 618 and 614, that allow objects in
30 the document view to be moved on the screen. The document view 622 may be further sub-divided into child views 602, 610 and 620 which may also overlap each other (and need not be rectangular). At any given time usually only one of the child views 602, 610 and 620 is active and only one view has input "focus". Only the view which has input focus responds to input actions and commands
35 from the input devices such as the mouse and the keyboard.

The toolbar/palette view 616 usually contains a number of iconic images, such as icons 608 and 612, which are used as a convenient way to initiate certain, often-used programs or subroutines. For example, icon 608 may be selected to

initiate a drawing routine which draws a box on the screen, whereas icon 612 might initiate a spelling checker program. The operation of such toolbars and palettes is generally well-known and will not be described further herein.

The displayed controls are generally selected by means of a mouse or
5 other input device. The mouse controls a cursor that is drawn on the screen by the operating system. When the cursor is positioned over the graphic image to be selected, a button is activated on the mouse causing the view system to respond.

Although the controls discussed above generally cannot be moved or
10 resized by the application thread, the content view and child views are usually totally under control of the application thread. When an application thread has several views, or several application threads, which are running simultaneously, and displaying information views, changes in the size or the position of one view will change the displayed or visible areas of views which
15 are "under" the changed view. Figures 7A and 7B illustrate how a manipulation of one view associated with an application can change the visible areas of other views that are associated with the same application and inside the same window.

In particular, Figure 7A shows three views located inside a window. The
20 views overlap - view 700 is in the background, view 702 is in front of view 700 and view 704 is in front of view 702. As shown in Figure 7A, view 704 obscures portions of views 702 and 700. Since each of views 700, 702 and 704 can be independently moved and resized, it is possible when the foremost views 702 or 704 are moved or resized, areas in the overlapped views can be uncovered or
25 covered and thereby change the visual appearance of these views. However, due to the overlapped appearance of the views, a change to a selected view only affects view behind the selected view. For example, a change to view 704 can affect views 702 and 700, but a change to view 700 cannot affect views 702 or 704 since these latter views overlap and obscure portions of view 700.

Figure 7B indicates the effect of a resizing of the front view 704 in Figure
30 7A. In particular, Figure 7B illustrates three views 712, 714 and 706, which correspond to views 704, 702 and 700 in Figure 7A, respectively. However, in Figure 7B, view 712 has been resized and, in particular, reduced in size from the original size view 704. The reduction in the size of view 712 exposes an area
35 (illustrated as shaded area) of view 710 that was previously totally covered by view 712. Similarly, the shaded portion 708 of view 706 is also uncovered. In accordance with normal view system operation, only visible portions of views are painted. Accordingly, areas 708 and 710 must be redrawn or repainted as

they have now become visible areas. These areas in need of redrawing are said to be "invalid" areas or "update" areas. This redrawing is accomplished by a coordination between the view system's background updating framework and the application thread which "owns" the invalid area as previously described.

5 Specifically, the view system computes the new visible area of each changed view and all views that lie behind the changed view. The view system then sends an "update request" to each view associated with a changed area indicating to the view that part of its visible area which must be redrawn. Each view, in turn, will proceed to update the visible area by directly writing into the
10 screen buffer.

The process of repainting a new visible area is shown in detail in the flowchart illustrated in Figure 8. In particular, the repainting routine starts in step 800 and proceeds to step 802 where the background updating framework receives an update request from the view system. An update request might be
15 generated when the view system resizes a view as illustrated in Figures 7A and 7B or in response to a newly uncover window area. In response, the updating framework acquires the view system monitor in function block 804, retrieves all the pending updates and combines them into a single update object as represented in step 806. A time stamp is then checked to see if the view system
20 display buffer is out of date, step 808. If the buffer is up-to-date, then it is copied to the display as in step 810. If the time stamp is out of date, then the view system display buffer must be recached. In step 812, a list of the views which overlap the update area is generated. At this point, the view system monitor is released 814. The background updating framework then iterates through each
25 view in the update list 816 and determines if the view has as associated application thread for processing updates 818.

If the view does not have a background updating thread of its own, then the view system background updating mechanism, itself, tells the view to redraw the invalid area into the view system display buffer 820. If the view
30 does have its own updating thread, then the invalid area within that view is handed off to the view's update thread 822 (which will cause the buffer to be updated and copied to the display at a later time). At this point, the up-to-date view system display buffer is copied to the display device buffer 824, the portion of the update associated with that view is marked as completed by the
35 background updating framework 826, and the thread goes back to waiting for further updates 828.

Also as previously mentioned, a view object can interact with the view hierarchy object to provide various view hierarchy management functions,

such as creating a new view and adding it to the hierarchy. An illustrative routine used by the application developer to create a new view is shown in detail in the flowchart of Figure 9. The routine starts in step 900 and proceeds to step 902 in which a new child view is constructed.

5 After the view is constructed and initialized in step 902, the routine proceeds to step 904 in which a "modifier" object is created to support and provide modifications to the view hierarchy. In step 908 the "Add" method on the modifier object is called and the newly constructed child view is added to the view hierarchy. The modifier object may then be thrown away as in step
10 910. The routine finishes in step 914.

As previously mentioned, the view system provides a framework that extends the intrinsic abilities of the C++ language to support construction and initialization of C++ objects. The steps involved in using the initialization and finalization support provided by the view system are detailed in the flow chart
15 Figure 10. The routine starts in step 1000 and the view object in question is constructed in step 1002 (using the standard C++ "constructor" facility.) After the construction of the object is complete, the "Initialize" method of the view object is called automatically by the run time system as illustrated in step 1005. A limitation of the C++ construction facilities makes it impossible to make
20 virtual method calls from within the constructor of an object. The "Initialize" method provides a mechanism for calling virtual methods on an object after the construction is complete, but before any other methods are called on the object. This is particularly useful for subclasses which need to modify base class behavior at construction or initialization time. Once the user quits the running
25 application, or some other event causes the view hierarchy (or even a single view) to be in need of destruction, a similar mechanism is used to insure that a "Finalize" method is called by the run time system after the last method call, but immediately before the object's destructor is run. This mechanism makes it possible to override behavior in the base class related to finalizing or destroying
30 an object. Step 1006 illustrates the removal of the view from the hierarchy, the run time system calling the view's "Finalize" method, step 1007, the destructor of the view being called, step 1008. The operation finishes in step 1010.

Figure 11 is a flow chart of an illustrative routine used by the view system to synchronize individual views with any changes in the view system
35 that might affect them. The routine starts in step 1100. The view monitor is acquired in step 1102. This monitor lock prevents further changes from happening in the view system for the duration of the routine. A time stamp within the view object is compared with the view system time "clock" in step

1104 to determine if the cached view state is now out of date, step 1106. If the cache is up-to-date, it is provided to the caller, step 1108. The view monitor is released 1122 and the routine completes. If the time stamp in step 1106 is out of date, then the view state is in need of recomputation. Step 1110 takes a copy of the parent view's state (which will go through a similar routine as outlined in figure 11 to provide its state to the requesting child view). The view object then adds in any local state information, for instance the distance the view is offset from its parent view and the difference between the area "owned" by the parent view and itself, step 1112. Once the up-to-date state has been computed, it is recached (step 114), along with the up-to-date time stamp, step 1116. The view monitor is released 1122 and the routine is completed, step 1124.

The schematic diagram in Figure 12 illustrates three different views that demonstrate the variety of view areas support by the view system framework. The enclosing window object 1200 is similar to the one described above in figure 6. The application developer-owned content view 1208 is contained within the window 1200. View 1206 is a child view of the content view which is non-rectangular, although of somewhat conventional shape. View 1202 is a non-rectilinear view which is a closed polygon and is also a child view of the content view 1208. The third child view 1204 in the content view 1208 is a non-contiguous view which has two separate, non-rectilinear portions. In this example, the child views do not overlap and a front to back ordering is not explicit. This is not meant to imply that ordering (and hence clipping) is not possible, only that it is not demonstrated in this figure.

The schematic diagram in Figure 13 illustrates two possibilities for coordinate spaces within any given view. The coordinate planes in any given view are usually invisible, but have been made visible in this figure for ease of explanation. The window 1200 is similar to that described in figure 6. The developer-owned content view 1202 is contained within the window 1200. Two child views of the content view are illustrated in this figure, view 1210 and view 1206. View 1206 is a view with the default coordinate system provided by the viewing framework. It consists of an "X" 1204 and "Y" 1208 axis that extend off into positive and negative infinity. By default, the "origin" is located in the top, left corner of the bounding rectangle of the view. The positive "X" direction extends to the right of the origin and the positive "Y" direction extends downward from the origin. In view 1206, the origin 1212 has been move slightly, but no other changes to the standard coordinate space have been made. In view 1210, a polar coordinate system is in operation. The origin 1216 of the view is located in the center of the view 1210. As in the Cartesian

coordinate system in view 1206, the polar coordinate system in view 1210 extends off into infinity as demonstrated by the arrow 1214. Both the illustrated coordinate systems are standard two dimensional coordinate planes. Three dimensional (3D) coordinate spaces are also possible, but not explicitly support
5 by the viewing framework. To use a 3D coordinate space, the developer would need to provide a 3D to 2D mapping for use by the viewing framework. In general, this is not a problem, because the graphics system implicitly supports 3D objects in a 2D coordinate plane.

The schematic diagram in Figure 14 illustrates the alignment and layout
10 mechanism provided by the viewing framework. The window 1400 is similar to the window described in figure 6. The developer-owned content view 1410 is contained within the window 1400. The views 1410, 1412, 1414 and 1416 are all child views of the content view. These views are objects provided by the viewing framework which are used explicitly for laying out other view objects.
15 Although it is not possible to actually see the layout views (they do not have any drawing capabilities), they are installed in the view hierarchy immediately above the views for which they hold layout attributes, and may be nested in the same fashion as other view objects. View 1402 is a child view of the layout view 1410. View 1404 is a child view of the layout view 1412. View 1406 is a
20 child view of the layout view 1414. View 1408 is a child view of the layout view 1416. Views 1410, 1412 and 1414 have attributes that cause them to be horizontally aligned at the center point. View 1416 and 1410 are also vertically aligned, both on the left side.

Figure 15 is a flow chart of an illustrative routine used by the view
25 system to arrange the views to produce the layout described in figure 14. The routine starts in step 1500. The view system queries a layout view in step 1502 for its horizontal and vertical attributes. In step 1560 the view system determines if the view is flexible in size. If it is flexible, the view may be resized to accommodate its layout wishes, step 1508. If the view is not flexible in size,
30 then it is not resized. Once any resizing is completed, the view may wish to be moved slightly to distribute white space between siblings, step 1510 and 1512. Once layout is completed, any display areas that have been changed are marked as in need of redrawing by the view system and an update request is posted to the background update framework. The routine is completed in step 1516.

35 The schematic diagram in Figure 16 illustrates a sample matrix transformation used within a view object which is used to provide an effect on another view object. Window 1600 is similar in operation to the window described in figure 6. The content view 1608 is contained within the window

1600. View 1602 is a rectangular view which contains the text "Hello, World". View effect 1604 is a "magnifier" view which displays the view(s) which are "underneath" it magnified by one and a half times (150%). This is accomplished by making use of the matrix object associated with the magnifier view and the support provided by the view system and graphic system to hold and use multiple matrix objects. View 1604 could be moved (by using a matrix translation as well as the matrix scaling) by the user over the content view 1608 to "magnify" other portions of the content view.

The schematic diagram in Figure 17 illustrates the mechanism provided by the viewing framework to "group" separate windows into a single layer, so that they "move" relative to other windows as a single group. For instance, the user moves the mouse over window 1702 and clicks the mouse button to select that window, window 1704 and 1706 will immediately be brought to the front of the display along with window 1702. Selecting any of the three windows 1702, 1704 or 1706 will cause all three windows to be selected and brought to the front of the display.

Figure 18 is a flow chart of an illustrative routine used by the view system to enable the use of non-multitask-safe objects by a framework which supports multitasking, as does the view system framework. This mechanism works without modifications to or internal knowledge of the workings of the non-multitask-safe object. The routine starts at step 1800. A method call is made on the non-multitask-safe object which is "caught" and implemented by the multitask-safe object wrapper, step 1802. In step 1804, the wrapper object checks to see if there are any non-multitask-safe objects not currently in use from within an internal "pool" of objects kept by the wrapper. If no objects are currently available, one is construed, step 1806 (another possible implementation might be to block and wait for another object in the pool to become available). Once an available object is obtained, the desired operation is completed using the non-multitask-safe object, as presented in step 1810. When the operation is completed, the internal monitor is released, step 1812 and the routine is completed, step 1814.

Figure 19 is a flow chart of an illustrative routine used by the view system to interact with the input system and distribute a positional event to its destination view object. To simplify this flow chart, no special distribution rules are provided or used. The routine starts at step 1900. In step 1902, a positional event is handed to the view system from the input system which received the event from some device that generates positional events, for instance a standard mouse device. The view monitor lock is acquired in step 1904. The

view system then makes use of the view hierarchy and a set of extensible rule objects to determine in which view the point is contained. The view system checks to make sure the point is within the "first" (root) view, step 1906. The operating system would not have passed the positional event to this window if
5 the point was not within the window's bounds, so for the root view case, the test should always return true. (If for some reason it returns false, then there may be a "bug" in some portion of the operating system.) Once it is determined that the positional event is within the root view, the view system determines if the point is contained within a child view of the root view, step 1910. If the
10 point is not within a child view, then the "target" for the positional event is the current view. If the point is contained within a child view, then the view system determines which child view, step 1910 and 1914 and continues recursively until the correct, front-most view which contains the positional event has been found (back to 1906 etc.) Once the target has been found, the
15 view monitor is released, step 1916, and the computed target is returned to the input system, step 1918 which will handle all further processing of the event. The routine is completed in step 1920.

Figure 20A and 20B are flow charts of illustrative routines used by the background updating framework to collect, batch and process changes in the
20 display area of one or more view objects in the view system. In figure 20A, the routine starts at step 2000. When a change occurs in a view object which is currently installed in the view hierarchy, notification of that change is posted to the view system, step 2002. The view system acquires the view monitor, step 2204 and then checks to see if changes are currently being "batched" for the view
25 system, step 2006. (When the view system is in "batching" mode, any redrawing that may need to be done as a result of view system changes is collected by the view system and saved for later processing by the background updating framework. If batching is not currently in use, then update requests are immediately posted to the background updating framework and processed.
30 In figure 20A, if changes are being batched, then the change notification is added to the list of changes being saved for later dispatch, step 2010. If changes are not currently being batched, then the change notification is immediately posted to any objects which have registered interest in that change, step 2008. Once the change notification has been handled (either saved or distributed), the view
35 monitor is released, step 2012 and the routine is finished, step 2014. Figure 20B illustrates the routine used to process change notifications once batching has been turned off. The routine starts at step 2020. In step 2022, the batching condition is changed to allow all pending change notifications to be posted. The

view system acquires the view monitor, step 2024, posts all the saved change notifications, step 2026 and then releases the view monitor, step 2028. The routine is completed at step 2030.

5 Figure 21A is a schematic diagram and figure 21B is a flow chart of mechanism used by the view system to provide read-only and read-write access to the view hierarchy. Figure 21A illustrates that iterator objects 2100 have access to only read-only methods 2104 of the view hierarchy 2108. Modifier objects 2102 have access to both read-only methods 2104 of the view hierarchy object 2108 as well as read-write methods 2106 (methods which may modify the
10 hierarchy) of the view hierarchy object 2108. Figure 21B demonstrates the steps performed in a sample operation of a iterator object (which is being used to compute the total area of all the views child views). The routine starts in step 2150. The view object constructs an iterator object in step 2152. Using the iterator object, the view retrieves the area for each of its direct child views and
15 adds it to the total area, step 2154. Once this operation is complete, the iterator is destructed, step 2158 and the routine is completed, step 2160. No monitors (or semaphores) are needed due to the fact that the iterator (and modifier) objects acquire and release the necessary monitor(s) from within their internal implementation.

20 While the invention is described in terms of preferred embodiments in a specific system environment, those skilled in the art will recognize that the invention can be practiced, with modification, in other and different hardware and software environments within the spirit and scope of the appended claims.

CLAIMS

Having thus described our invention, what we claim as new, and desire to secure by Letters Patent is:

- 1 1. An apparatus for initializing a system environment, comprising:
 - 2 (a) a processor;
 - 3 (b) a storage attached to and under the control of the processor;
 - 4 (c) a object-oriented operating environment resident in the storage for
5 controlling the processor;
 - 6 (d) an object processor in the operating environment for constructing a first
7 object with an associated data area and method; and
 - 8 (e) a virtual base class object for invoking the object processor of the
9 operating environment to polymorphically initialize the first object.

- 1 2. The apparatus as recited in claim 1, including means within the object-
2 oriented operating environment for initializing a class identification
3 object with a unique identifier for identifying the first object.

- 1 3. The apparatus as recited in claim 2, including means within the class
2 identification object for identifying the first object with the unique
3 identifier.

- 1 4. The apparatus as recited in claim 1, including object processing means for
2 querying information to determine if the first object has been properly
3 initialized.

- 1 5. An apparatus for finalizing a system environment, comprising:
 - 2 (a) a processor;
 - 3 (b) a storage attached to and under the control of the processor;
 - 4 (c) a object-oriented operating environment resident in the storage for
5 controlling the processor;
 - 6 (d) an object processor in the operating environment for destructing a first
7 object with an associated data area and method; and
 - 8 (e) a virtual base class object for invoking the object processor of the
9 operating environment to polymorphically finalize the first object.

- 1 6. The apparatus as recited in claim 5, including means within the object-
2 oriented operating environment for initializing a class identification
3 object with a unique identifier for identifying the first object.
- 1 7. The apparatus as recited in claim 6, including means within the class
2 identification object for identifying the first object with the unique
3 identifier.
- 1 8. The apparatus as recited in claim 7, including object processing means for
2 querying information to determine if the first object has been properly
3 initialized.
- 1 9. A method for initializing a system environment with a processor, a
2 storage attached to and under the control of the processor and a object-
3 oriented operating environment resident in the storage for controlling
4 the processor, comprising the steps of:
5 (a) constructing a first object with an associated data area and method
6 utilizing the object processor in the operating environment; and
7 (b) invoking the object processor of the operating environment to
8 polymorphically initialize the first object utilizing a virtual base class
9 object.
- 1 10. The apparatus as recited in claim 9, including the step of initializing a
2 class identification object with a unique identifier for identifying the first
3 object.
- 1 11. The apparatus as recited in claim 10, including the step of identifying the
2 first object with the unique identifier located within the class
3 identification object.
- 1 12. The apparatus as recited in claim 11, including the step of determining if
2 the first object has been properly initialized.
- 1 13. A method for finalizing a system environment with a processor, a
2 storage attached to and under the control of the processor and a object-
3 oriented operating environment resident in the storage for controlling
4 the processor, comprising the steps of:
5 (a) destructing a first object with an associated data area and method
6 utilizing the object processor in the operating environment; and

- 7 (b) invoking the object processor of the operating environment to
8 polymorphically finalize the first object utilizing a virtual base class
9 object.
- 1 14. The method as recited in claim 13, including the step of initializing a
2 class identification object with a unique identifier for identifying the first
3 object.
- 1 15. The method as recited in claim 14, including identifying the first object
2 with the unique identifier utilizing the class identification object.
- 1 16. The method as recited in claim 15, including determining if the first
2 object has been properly initialized.




FIG. 1 (Prior Art)

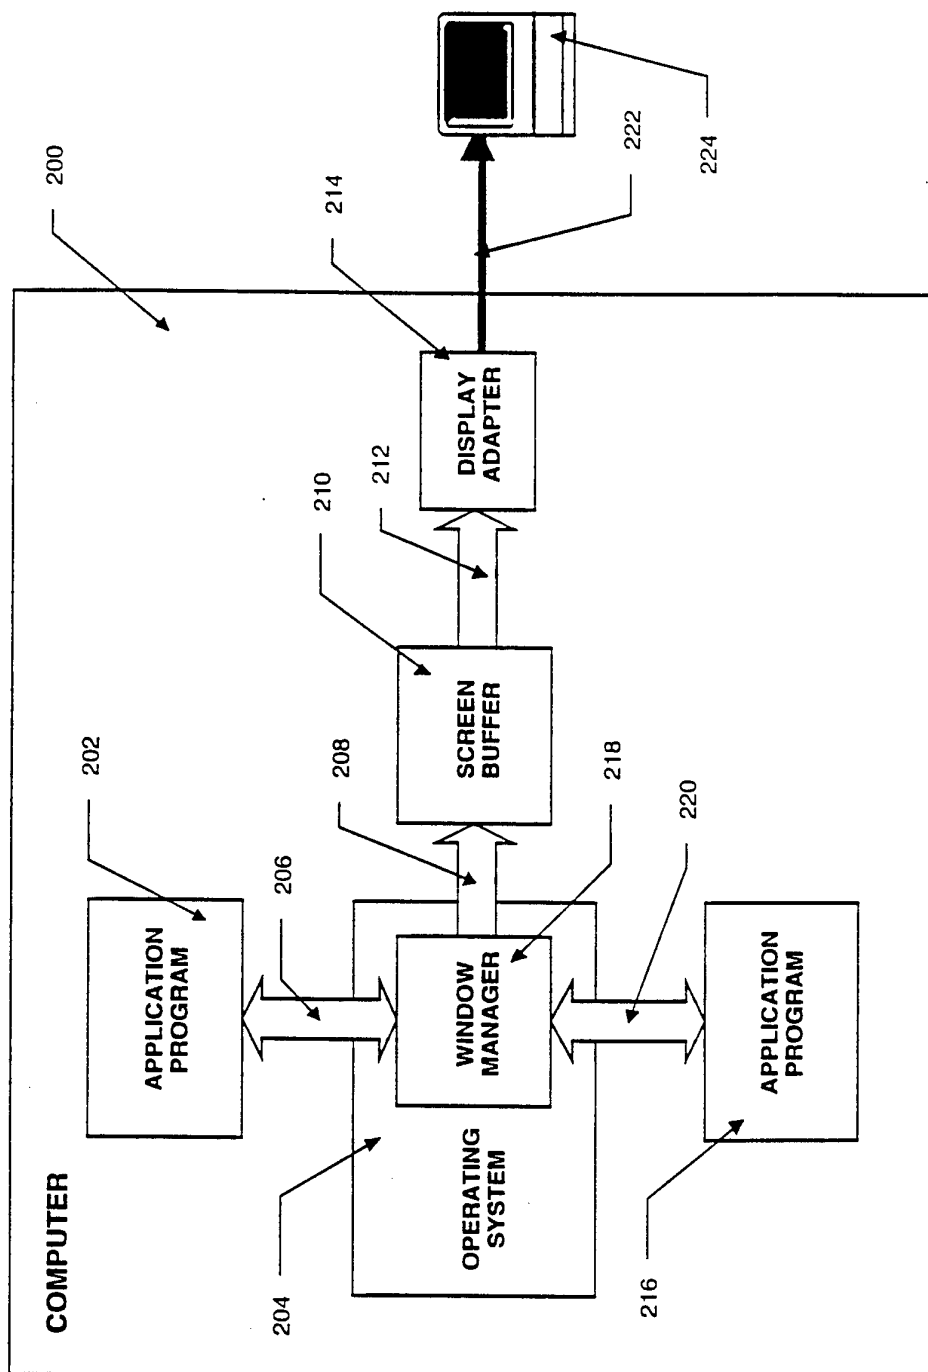


FIG. 2 (Prior Art)




FIG. 3




FIG. 4

5 / 2 2




FIG. 5

6/22




FIG. 6




FIG. 7A




FIG. 7B

8 / 2 2




FIG. 8A

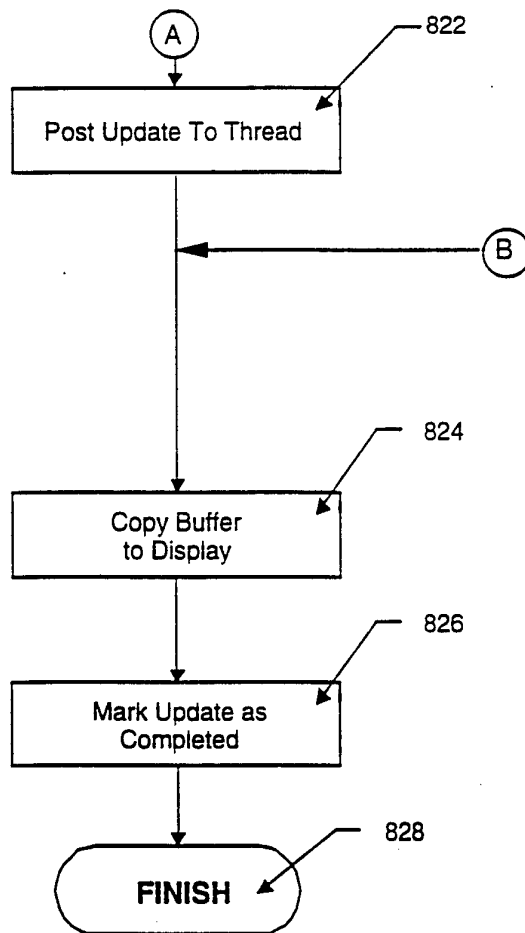


FIG. 8B

1 0 / 2 2

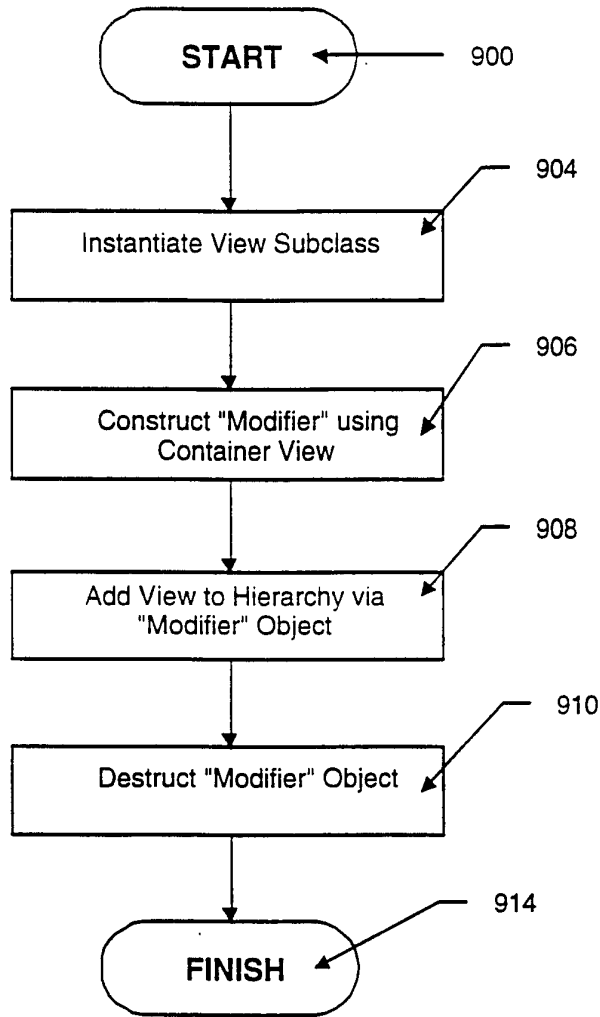


FIG. 9

11/22

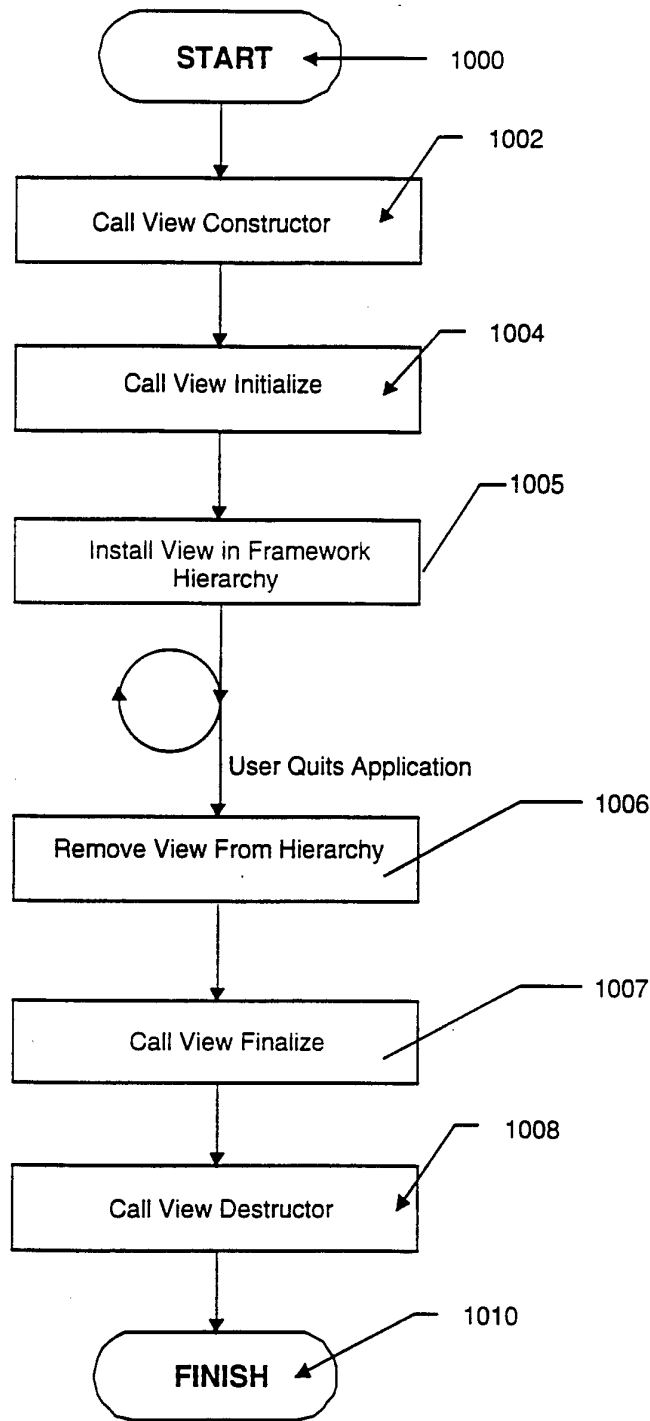


FIG. 10

1 2 / 2 2

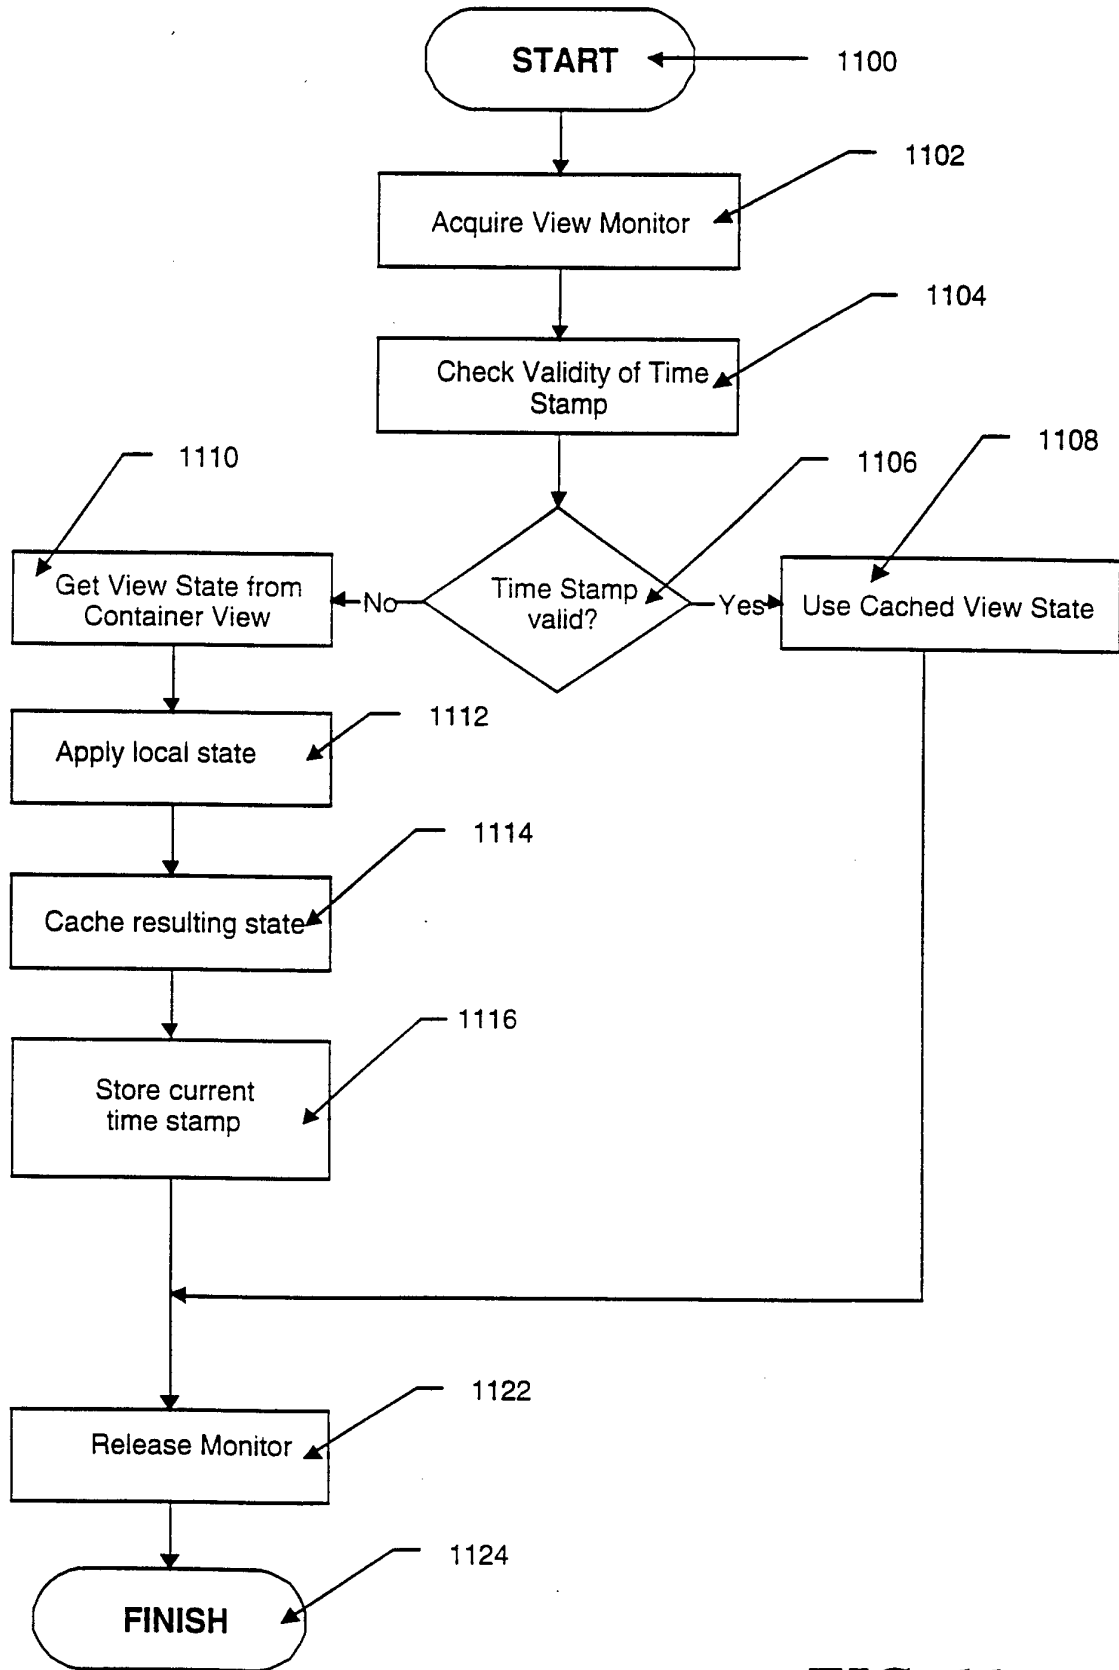


FIG. 11




FIG. 12




FIG. 13




FIG. 14

1410

1400

16/22

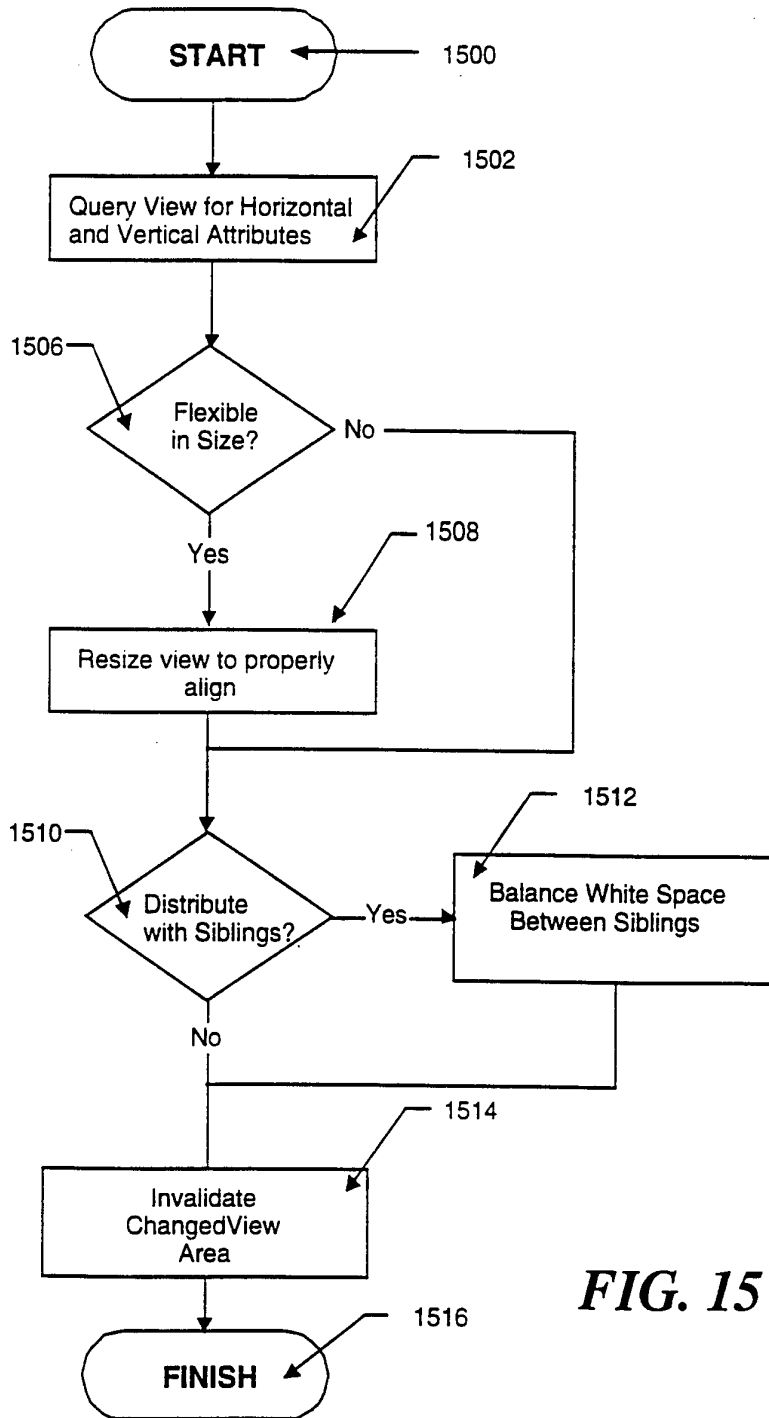


FIG. 15




FIG. 16

18/22




FIG. 17

19/22




FIG. 18

20/22

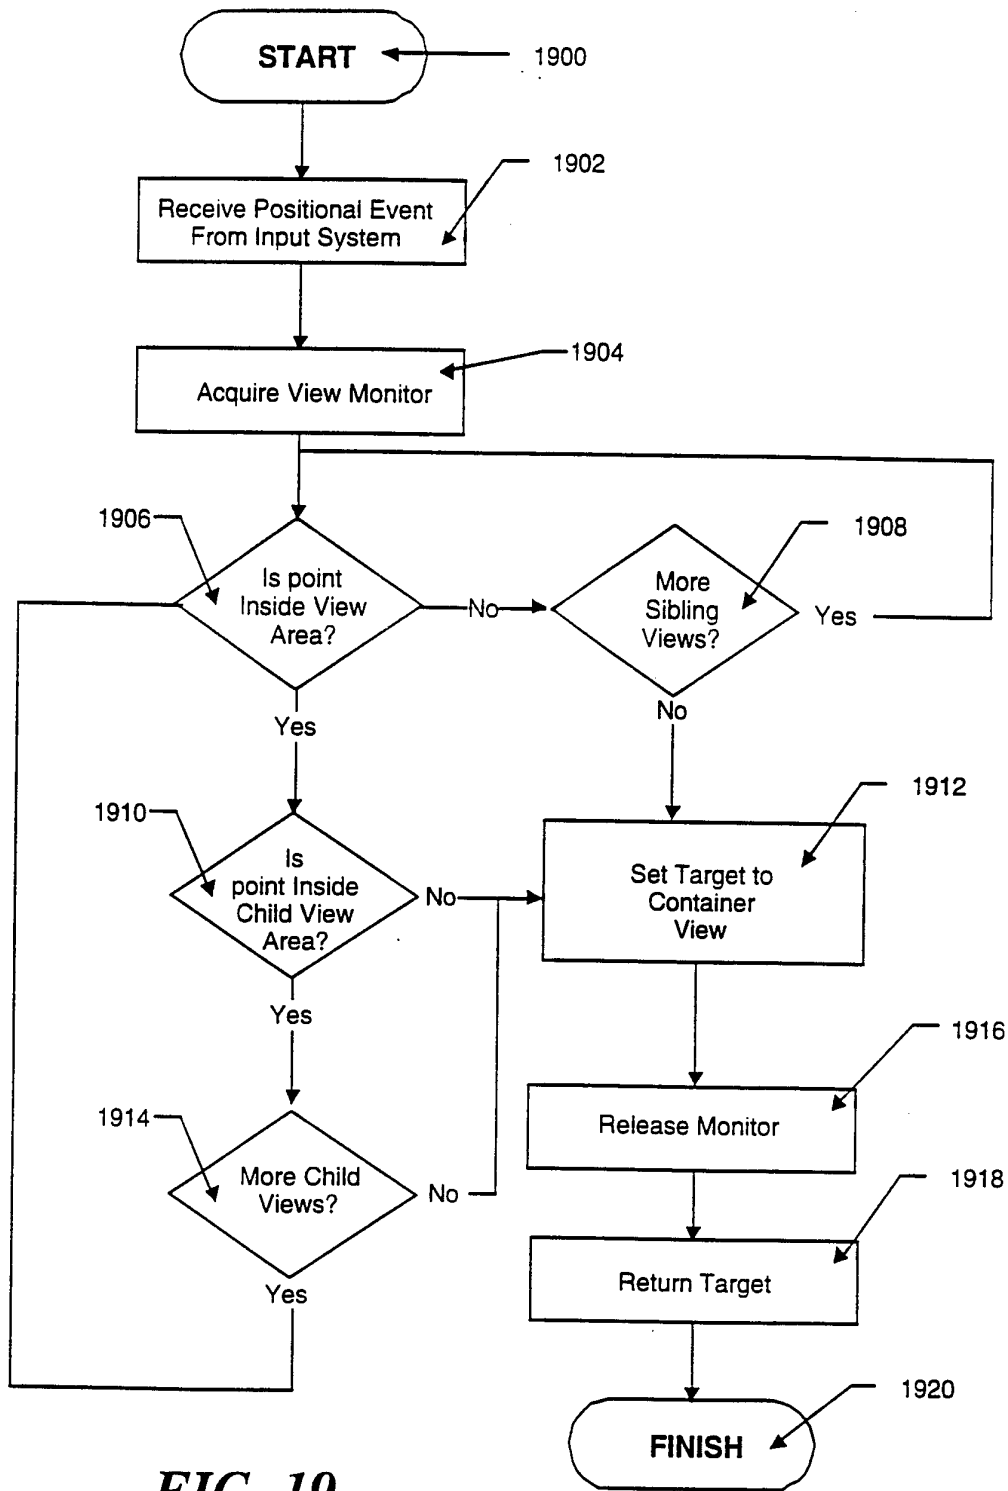


FIG. 19

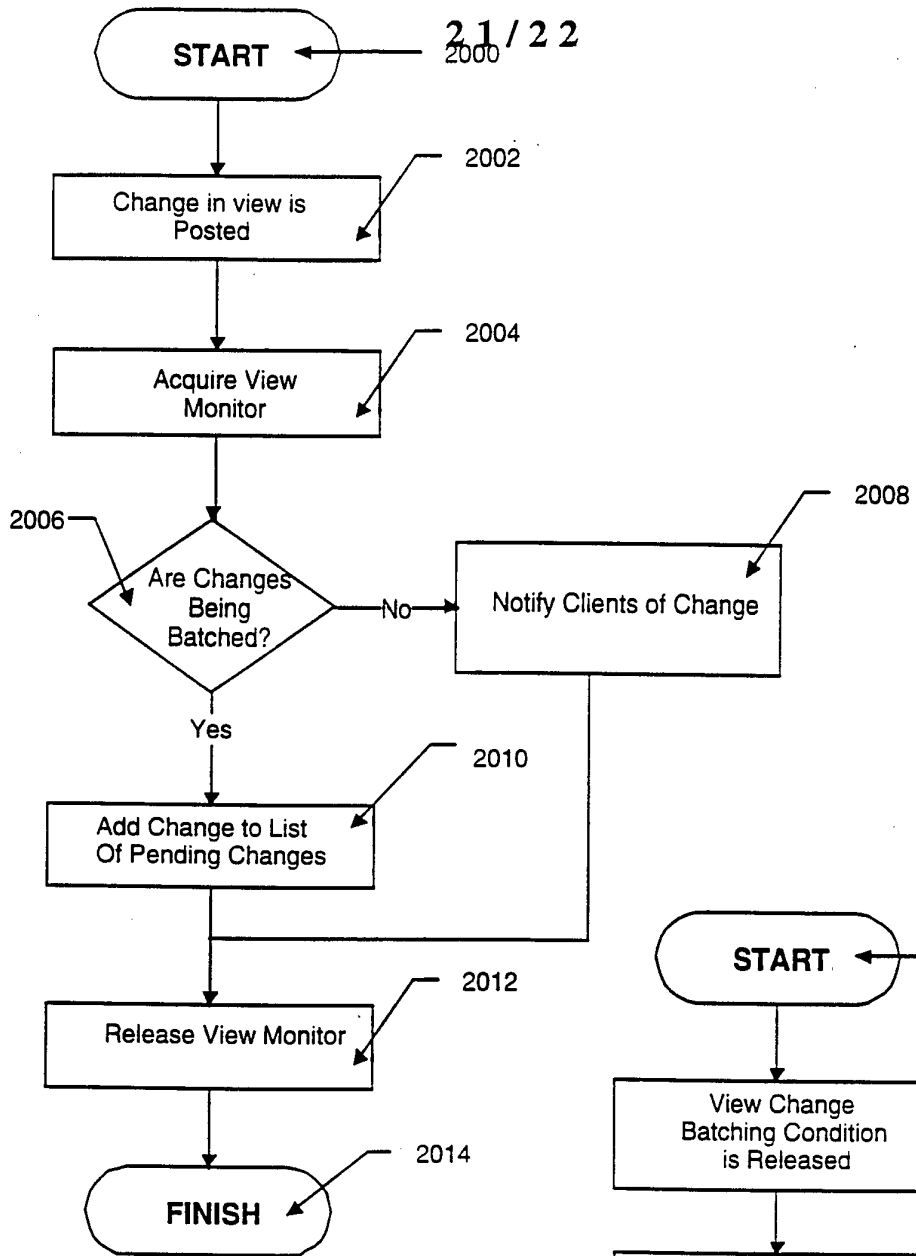


FIG. 20A




FIG. 20B

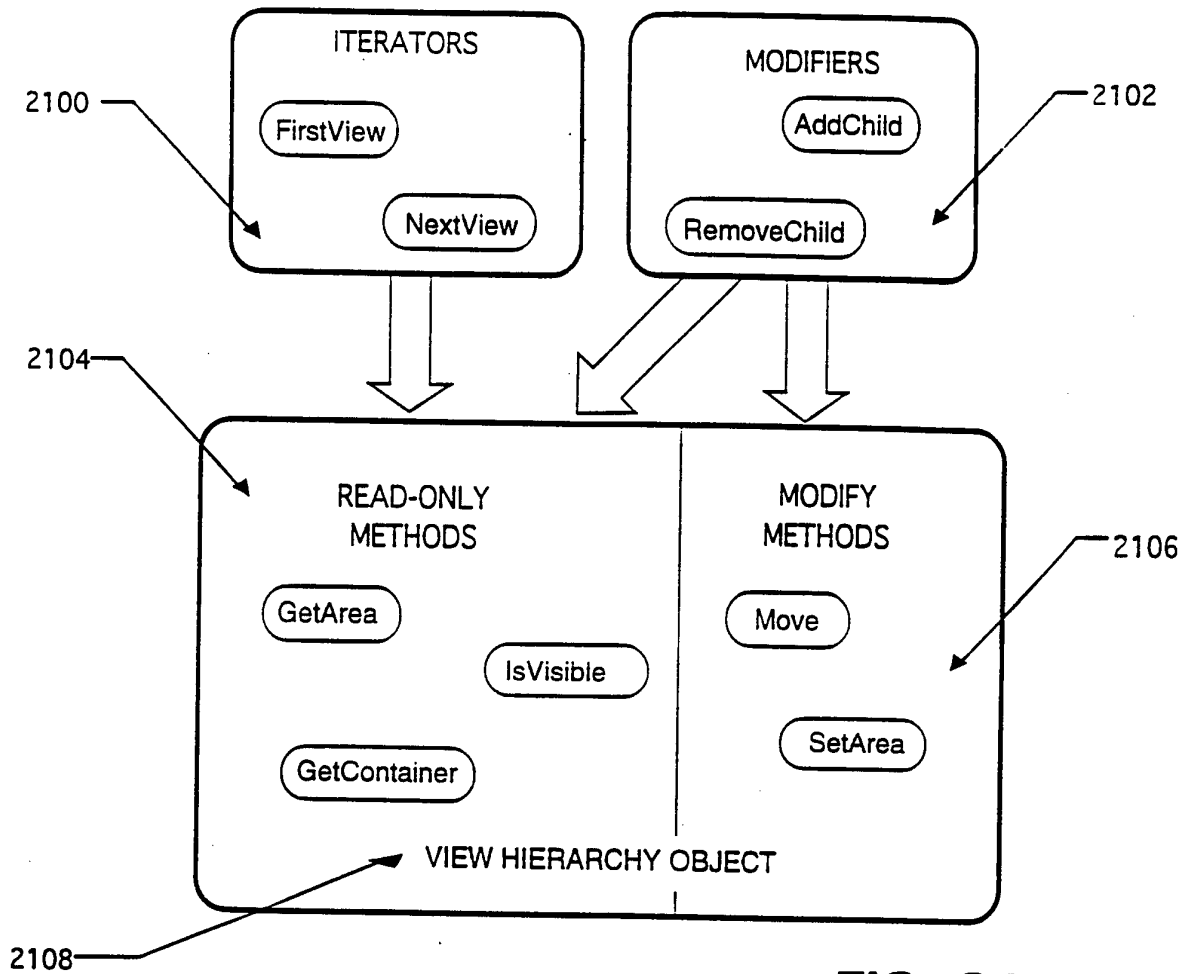


FIG. 21A




FIG. 21B

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 94/10195

A. CLASSIFICATION OF SUBJECT MATTER IPC 6 G06F9/44				
According to International Patent Classification (IPC) or to both national classification and IPC				
B. FIELDS SEARCHED				
Minimum documentation searched (classification system followed by classification symbols) IPC 6 G06F G09G				
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched				
Electronic data base consulted during the international search (name of data base and, where practical, search terms used)				
C. DOCUMENTS CONSIDERED TO BE RELEVANT				
Category ^o	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.		
A	EP,A,0 476 635 (HEWLETT-PACKARD COMPANY) 25 March 1992 see page 2, line 6 - page 2, line 21 ---	1		
A	EP,A,0 569 861 (MICROSOFT CORPORATION) 18 November 1993 see abstract -----	1		
<input type="checkbox"/> Further documents are listed in the continuation of box C.				
<input checked="" type="checkbox"/> Patent family members are listed in annex.				
^o Special categories of cited documents :				
<table style="width: 100%; border: none;"> <tr> <td style="width: 50%; vertical-align: top; padding: 5px;"> "A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier document but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed </td> <td style="width: 50%; vertical-align: top; padding: 5px;"> "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art. "&" document member of the same patent family </td> </tr> </table>			"A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier document but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art. "&" document member of the same patent family
"A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier document but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art. "&" document member of the same patent family			
Date of the actual completion of the international search <p style="text-align: center; font-size: 1.2em;">9 March 1995</p>	Date of mailing of the international search report <p style="text-align: center; font-size: 1.2em;">14.03.95</p>			
Name and mailing address of the ISA European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+ 31-70) 340-2040, Tx. 31 651 epo nl, Fax (+ 31-70) 340-3016	Authorized officer <p style="text-align: center; font-size: 1.2em;">Van Roost, L</p>			

1

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US 94/10195

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP-A-476635	25-03-92	US-A- 5361351 JP-A- 4247536	01-11-94 03-09-92
EP-A-569861	18-11-93	US-A- 5327562 JP-A- 6103085	05-07-94 15-04-94