



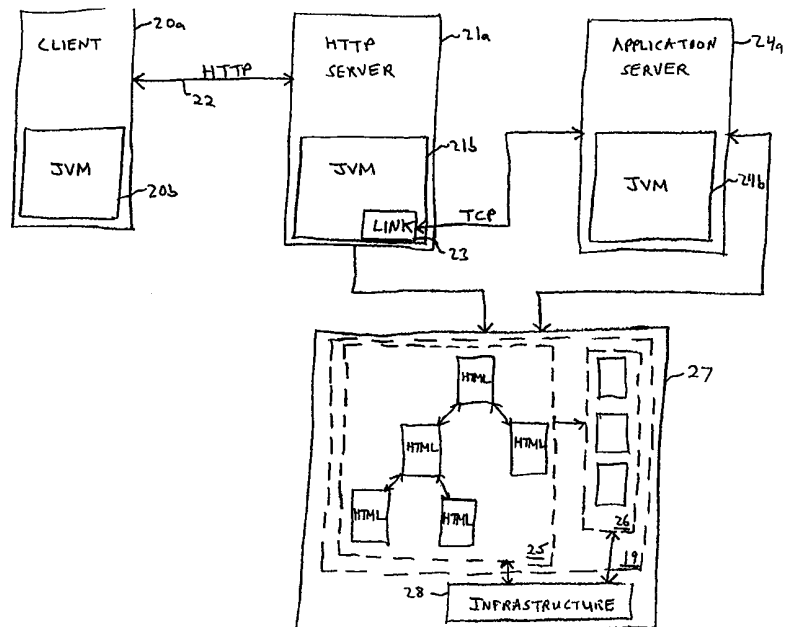
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification ⁶ : G06F 9/42, 15/16, 17/30</p>	<p>A1</p>	<p>(11) International Publication Number: WO 99/27440 (43) International Publication Date: 3 June 1999 (03.06.99)</p>
<p>(21) International Application Number: PCT/US98/25275 (22) International Filing Date: 25 November 1998 (25.11.98) (30) Priority Data: 08/978,349 25 November 1997 (25.11.97) US (71) Applicant (for all designated States except US): ALPHA BLOX CORPORATION [US/US]; 800 Maude Avenue, Mountain View, CA 94043-4020 (US). (72) Inventors; and (75) Inventors/Applicants (for US only): DALE, Geoffrey, W. [US/US]; 21090 Grenola Drive, Cupertino, CA 95014 (US). SWENSON, Eric, J. [US/US]; 745 Granite Ridge Drive, Santa Cruz, CA 95065 (US). SKOK, Michael, J. [GB/US]; 54 Serano Drive, Atherton, CA 94027 (US). STAVE, Matthew [US/US]; 3100 La Terrace Circle, San Jose, CA 95123 (US). POONEN, Sanjay, J. [IN/US]; 932 Providence Court, Cupertino, CA 95014 (US). (74) Agents: SALTER, James, H. et al.; Blakely, Sokoloff, Taylor & Zafman LLP, 7th floor, 12400 Wilshire Boulevard, Los Angeles, CA 90025 (US).</p>		<p>(81) Designated States: AL, AM, AT, AT (Utility model), AU (Petty patent), AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, CZ (Utility model), DE, DE (Utility model), DK, DK (Utility model), EE, EE (Utility model), ES, FI, FI (Utility model), GB, GD, GE, GH, GM, HR, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SK (Utility model), SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, 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 <i>With international search report.</i></p>

(54) Title: TIER-NEUTRAL DEVELOPMENT OF HYPERTEXT BASED APPLICATIONS

(57) Abstract

A tier-neutral development mechanism for network-based applications is provided. An application created using the mechanism includes a plurality of hypertext pages (item 25), at least some of which incorporate executable components. The application is invoked by a hypertext request for a page from a browser running on a client tier (20a). An application server (24a) responds to the request by retrieving the requested page and assigning any components incorporated therein to the proper tier for execution. The mechanism provides a single model by which any executable component can be specified by an application developer for execution on any tier on the network or made subject to an automatic, dynamic tier assignment by the application server (24a). Components of a given application can be distributed across, and specified for execution on, three or more different tiers and moved from tier to tier. An application developer can use a conventional hypertext editor to integrate selected components into extended hypertext pages to create an application.



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			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

Tier Neutral Development of Hypertext Based Applications

FIELD OF THE INVENTION

The present invention pertains to the field of network-based software. More particularly, the present invention relates to software applications that are designed for use on the Internet.

BACKGROUND OF THE INVENTION

The rapid growth in the use of the Internet and the World Wide Web ("the Web") has sparked a dramatic increase in the development of Internet-based software applications. Developments in Internet software have arisen in many different areas ranging from on-line shopping and banking to games and entertainment. Web-based tools such as Java, JavaBeans, and ActiveX have helped to partially bridge the gap between the simple, text-only Web pages of the early Internet and sophisticated software applications. In particular, these tools allow Web designers to include small, special-purpose programs (sometimes referred to as "applets", "Beans", or more generally, "components") within Web pages written in a hypertext language, such as the Hypertext Mark-up Language (HTML). Thus, Web pages have been transformed from passive, text-only displayable documents to dynamic documents that can provide moving graphics, video, sounds, and interactive capability.

While progress is being made in the development of Internet applications, application development tools for the Internet still have a number of shortcomings. For example, conventional Web-oriented programming models such as Java, JavaBeans, and ActiveX are not "tier-

neutral" with respect to the environments in which they run. That is, the services upon which such models rely tend to be very specific to the tier (e.g., client or server), operating system, and in some cases, hardware. For example, while Java and JavaBeans do not have operating system and hardware dependencies as ActiveX controls do, they must exist in an environment or infrastructure that provides the necessary services (e.g., communications, access to environment, parameterization, connectivity to external services, etc.) to operate.

The surrounding environment or infrastructure in prior art application development systems tend to be tier-specific. The tools, mechanisms, and services available on client tiers, for example, are quite different from those available on server tiers. In heterogeneous hardware environments, the server tiers may also provide significantly different environments. Further, these systems do not allow components of an application to be moved from one tier to another, or to be distributed among multiple tiers. Such capability may be desirable for purposes of performance, security, or browser capability.

In addition, with conventional tools, interfaces between applets or other executable components must be custom-coded into the hypertext document. In particular, an application developer specifically "wires" (i.e., writes code for) the connections between components, so that the components can communicate with each other. This constraint makes it difficult for inexperienced programmers to create Web applications and causes application development to be a laborious, time-consuming process.

A common approach to creating application development tools for the Web has been to retrofit existing tools with Web "front ends" (user interfaces). As a result, such tools are not ideally-suited to the

development of Web applications. Further, retrofitted tools are not likely to adapt well to future changes in the platforms, users, and types of usage of the Internet.

Hence, it would be desirable to have Internet applications and application development mechanisms which overcome the above-mentioned disadvantages. In particular, what is needed is an application development mechanism that is specifically designed for development of Internet applications, especially Web-based applications. What is further needed is a single development mechanism that is based on a single, tier-neutral model, using which application developers can create application components for execution on any tier on a network, distribute application components across multiple tiers on a network, or move application components from tier to tier. In addition, what is needed is an application development mechanism which allows Internet applications to be created faster and more easily by people who have little or no programming experience.

SUMMARY OF THE INVENTION

The present invention includes a mechanism for creating a software application for execution on a network that has multiple tiers. The mechanism comprises means for specifying application components, which allows any specified application component to be assigned to execute on any of the tiers of the network. The mechanism further comprises means for enabling the application components to be associated with one or more hypertext-based pages, which allows the components to be executed in response to requests for the hypertext-based pages. Other features of the present invention will be apparent

from the accompanying drawings and from the detailed description which follows.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

Figure 1 illustrates a network arrangement by which a Web server provides a client computer with access to HTML pages.

Figure 2 illustrates a computer system.

Figure 3 illustrates a network arrangement by which a server computer provides a client computer with access to a Web-based application in accordance with the present invention.

Figure 4A illustrates software application components in communication with an infrastructure.

Figure 4B illustrates a software application distributed across multiple tiers.

Figure 5 is a flow diagram illustrating a routine by which a software application is invoked.

Figure 6 is a flow diagram illustrating a routine by which a component is instantiated and registered.

Figure 7 illustrates in greater detail the components illustrated in Figure 4.

Figure 8A is a flow diagram illustrating in greater detail certain steps of the routine of Figure 6.

Figure 8B is a diagram illustrating three executable components of an application being executed on three separate tiers.

Figure 9 is a flow diagram illustrating a routine by which an interface between two components is established.

Figure 10 is a diagram illustrating multiple executable components of an application being executed on multiple tiers.

Figure 11 illustrates a display screen showing a graphical user interface for assembling a software application.

Figure 12 illustrates a display screen showing a graphical user interface, using which an end user can assemble and execute a customized application.

DETAILED DESCRIPTION

A tier-neutral development mechanism for Web-based applications is described. In the following description, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be evident, however, to one skilled in the art that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to facilitate description of the present invention.

The present invention includes a "tier-neutral" mechanism for application assembly, using which developers with little or no programming experience can easily create Web based applications. In other words, using a single application development mechanism in accordance with the present invention, Web developers can create applications to run on any tier on a network or to run across multiple tiers. In this description, the term "tier" is used to distinguish between a

client and a server on a network and to distinguish between multiple servers.

The present invention may be carried out in one or more computer systems, each in response to its central processing unit (CPU) executing sequences of instructions contained in memory. That is, execution of the sequences of instructions contained in memory causes the computer's CPU to perform steps to execute part or all of a Web-based application. As will be apparent from the description which follows, the instructions may be loaded into the computer's memory from a persistent store, such as a mass storage device, and/or from one or more of the other computer systems (collectively referred to as the "host computer system") over the Internet. For example, the host computer system may transmit a sequence of instructions to the receiving ("target") computer system in response to a message transmitted to the host computer system over a network by the target computer system. As the target computer system receives the instructions via a network connection, such as a conventional telephone modem, Ethernet adapter, Integrated Services Digital Network (ISDN) adapter, or the like, the target computer system stores the instructions in memory. The target computer system may store the instructions for later execution or execute the instructions as they arrive over the network connection. The host computer system may also execute various aspects of the Web-based application.

In some cases, instructions which embody the present invention may be directly supported by the CPU of the executing computer. Consequently, execution of the instructions may be performed directly by that computer's CPU. In other cases, the instructions may not be directly executable by the computer's CPU; under these circumstances, the instructions may be executed by causing the CPU to execute an

interpreter that interprets the instructions, or by causing the CPU to execute instructions which convert the received instructions to instructions which can be directly executed by the CPU.

In alternative embodiments, hardwired circuitry may be used in place of, or in combination with, software instructions to implement the present invention. Thus, the present invention is not limited to any specific combination of hardware circuitry and software, nor to any particular source for the instructions executed by a computer system.

Figure 1 illustrates an arrangement of processing systems by which a user can browse the World Wide Web ("the Web"). A client computer system 1 has a connection 3 via the Internet to a server computer system 2. The client system 1 executes browser software which, in response to a user command, sends a request for a Hypertext Mark-up Language (HTML) page to the server 2 using Hypertext Transport Protocol (HTTP). The requested HTML page is one of a number of HTML pages 4 that are stored in a storage facility 5. The storage facility 5 may or may not be physically located on the same computer system as the server 2. The HTML pages 4 are functionally interconnected by various hypertext links and may be logically arranged in a tree hierarchy, as shown. In response to the browser's request, the server 2 retrieves the requested page and provides it to the client 1, which then displays the page to the user.

Figure 2 is a simplified block diagram of the architecture of a computer system, such as a conventional personal computer (PC). Client system 1, server system 2, and any other computer system mentioned in this description may have the architecture of Figure 2 or a similar architecture, except as stated otherwise. Numerous variations upon the overall configuration illustrated in Figure 2 are possible, however, within the scope of the present invention. The computer system of Figure 2

includes a CPU 10, random access memory (RAM) 11, read-only memory (ROM) 12, and a mass storage device 13, all coupled to a bus 18. The bus 18 may actually include one or more physical buses interconnected by one or more adapters, bridges and/or controllers. Also connected to the bus 18 are a display device 14, a keyboard 15, a pointing device 16, and a communication device 17. The communication device 17 provides an interface between the illustrated computer system and the network connection 3 for purposes of communicating with other computer systems.

The communication device 17 may be any suitable device for providing digital communication over the network connection 3, such as a conventional Ethernet or other local area network (LAN) adapter, telephone modem, ISDN adapter, Digital Subscriber Line (xDSL) adapter, cable television modem, or any other suitable device. The mass storage device 13 may include any suitable non-volatile medium for storing large amounts of data, such as a magnetic disk, a magneto-optical (MO) disk, a CD-ROM, CD-R, CD-RW, or the like. The pointing device 16 may be any suitable device for allowing a user to position a pointer displayed on the display device 14, such as a mouse, trackball, touchpad, or the like. Numerous variations upon the aforementioned components are possible within the scope of the present invention.

The present invention includes a "tier-neutral" model by which Web developers can easily create Web based applications. In other words, using a single application assembly mechanism of the present invention, Web developers can create applications for any tier and, if desired, distribute applications across multiple tiers. Further, the tool allows applications to be created quickly and easily by someone with little or no programming experience.

Using the above-mentioned development mechanism, Web based applications are created from conventional HTML pages in combination with certain executable "components" and a program infrastructure, as will be described below. The specific functions of the executable components depend upon the nature of the application but may include, for example, charting or spreadsheet functions, spell checking, and various user interface functions. The components can be written in Java programming language and implemented as Java classes. In such embodiments, the components can be incorporated into HTML pages using applet tags. Thus, an application designed according to the present invention includes one or more conventional HTML pages with one or more of these components. The components are dynamically downloaded as needed when the user browses to a corresponding HTML page, and the entire application is maintained on the server side.

Figure 3 illustrates one arrangement by which the user of a client computer system 20a can access and execute a Web-based application 19 created in accordance with the present invention. In the description which follows, the present invention is described using the example of an application which resides and executes on a wide area network (WAN), i.e., the Internet. Note, however, that the present invention can also be used to assemble applications which reside and execute on a LAN, such as in an Intranet environment. Referring then to Figure 3, the client computer system 20a executes browser software and preferably includes a Java Virtual Machine (JVM) 20b. The client 20a executes browser software and communicates over the Internet with a conventional Web server (hereinafter the "HTTP server") 21a via a network connection 22 using HTTP protocol. Network connection 22 may be a conventional telephone link, ISDN link, xDSL link, cable television link, Ethernet link, or any

other suitable communication link. The HTTP server 21a may include its own JVM 21b, as shown, although that is not necessary for purposes of the present invention. The HTTP server 21a may comprise any conventional Web server software, such as Netscape Enterprise Server, Microsoft Internet Information Server, or Apache.

In accordance with the present invention, the network arrangement also includes a second server 24a, which shall be referred to as the "application server" 24a. As will be described below, the application server 24a performs certain functions associated with the present invention, such as assigning executable application components to appropriate tiers. The HTTP server 21a and the application server 24a may or may not be physically located on the same computer system. The "application server" includes a JVM 24b. The JVM 21b of the HTTP server 21a includes a link component 23, which provides a communication interface between the HTTP server 21a and the application server 24a using Transport Control Protocol (TCP). The link provides a mechanism by which the HTTP server 21a can delegate HTTP requests to the application server 24a. If the HTTP server 21a does not include its own JVM 21b, this link mechanism can be provided using, for example, CGI (Common Gateway Interface), NSAPI (Netscape Server Application Program Interface), or Microsoft ISAPI (Internet Server API).

The HTTP server 21a and the application server 24a each have access to the contents of a storage facility 27. The storage facility 27 may or may not be physically located on the same computer system as the HTTP server 21a and/or the application server 24a. In fact, the storage facility 27 may be located essentially anywhere on the Internet and may be distributed across multiple computer systems. Further, the storage facility may be embodied as separate data repositories for the HTTP

server 21a and the application server 24a, or one storage facility may be a subset of the other. Thus, the storage facility 27 is described herein as a single logical and physical entity only for purposes of simplicity. The storage facility 27 may be embodied as one or more conventional storage devices.

The storage facility 27 stores a number of HTML pages 25 and a number of executable components 26 of the present invention. The HTML pages 25 in combination with the components 26 form a Web-based application 19 in accordance with the present invention. The HTML pages 25 may be logically arranged in a tree hierarchy, as shown. Generally, the application 19 is invoked in response to a request from the client 20a for an HTML page that includes one or more references ("tags") to one or more of the components 26. Execution of the application 19 is supported by a programming infrastructure 28, the purpose and functions of which are described below.

The components 26 of the present invention are based on the conventional concept of "applets", i.e., small portions of executable code that can be incorporated (by reference) into HTML pages. The components 26 may be embodied as Java classes, as is assumed to be the case for purposes of this description. However, it will be recognized that the components 26 may alternatively be implemented in another object-oriented language. Thus, the present invention is not limited to a Java implementation. Each of the components 26 is incorporated into one or more of the HTML pages 25 by placing an applet tag corresponding to the component within the appropriate HTML page or pages.

The application components 26 have several distinguishing characteristics, however, which make them conducive to rapid and easy application assembly. In particular, the components have relatively large

granularity as well as fully exposed controls, which allow the components to be used by all levels of users from casual page designers to application developers. The components differ by an order of magnitude in granularity from traditional applets. Because the components tend to be inherently larger than traditional applets, they may encompass an entire business process or function. For example, a single component can be a charting component or a spreadsheet component.

The components 26 can execute on any tier on the network, e.g., only on the application server 24a, only on the client 20a, or on any available tier (e.g., one of several available servers). User interfaces of the components 26 are completely decoupled and can be rendered on different platforms using different languages. Further, the components 26 have features which allow connections between them to be automatically established at run time without any scripting or custom coding. All components have fully exposed properties (e.g., applet parameters), methods, and events that can be controlled through script, for example. Alternatively, properties can be set at design time using a development tool or text editor and become the initial state of a component. If the application developer is using any conventional page design tool, such as Netscape Visual JavaScript, Microsoft FrontPage, or Net Objects Fusion, the property style sheets can be presented in simple dialog boxes, and the exposed properties, methods, and events can be scripted through point-and-click dialogs.

As noted above, the components 26 can execute on any tier. More specifically, the system of the present invention allows the application developer to make an explicit choice of tier assignment for any component, if desired. Accordingly, each component can be explicitly specified to execute on the client tier, on the server tier, or on a specific

one of multiple server tiers. Alternatively, the application server 24a can automatically decide on which tier a component is to execute, such as when no tier specification is made for the component. This automatic tier assignment may be made based on, for example, available resources, load balancing criteria, hints, performance criteria, or other criteria. In that case, the application developer need not be concerned with the tier or tiers on which a component should execute. The ability to make an automatic, dynamic tier assignment, therefore, provides an advantage over other application development systems, in which the developer must explicit choose a tier.

Server-side components (i.e., components which execute on a server tier) can be specified to execute so that their output is rendered as HTML back into the referencing HTML page in place of the applet tag before the referencing page is provided to the client 20a. In that case, the client 20a receives only a conventional HTML page with no applet tags.

Another important feature of the present invention is that, in contrast with conventional Web-based development tools, functional connections between components of a Web application can be established to allow components to communicate with each other, without the need for custom coding by the application developer to specifically "wire" (specify) the connections. Thus, the tier-neutral model of the present invention allows the creation of components which effectively "snap together" automatically when the application is executed. This characteristic is in contrast with conventional programming models which require custom coding within hypertext pages to create connections between applets.

A simple example of how the application 19 executes is described now with reference to Figure 4A. For purposes of description, assume

that the components 26 have been assigned to execute on the client tier 20a, unless stated otherwise. However, note that many of the functions described below can also be performed in like manner by the application server 24a. In the example of Figure 4A, the application 19 allows a user of the client 20a to perform a query upon a database (not shown). The result set of the query is to be provided by a table. Consequently, the components 26 include a query component 31 for receiving user input specifying the query and a table component 30 for actually querying the database and for providing a result set in response to the query. Note that the table component 30 may also provide the result set to some form of output component (not shown) for displaying a chart or some other form of report to the user.

At this point, it should be noted that in this description, certain components, methods, and routines are described as "performing", "for performing", or "responsible for" a particular function, or other similar language. It will be recognized, however, that it is the execution of these components, methods, and routines by a processor (i.e., a CPU) which actually causes the function to occur.

To perform the query upon the database, a functional connection 32 is required between the query component 31 and the table component 30. With conventional programming models, such a connection would be established by custom coding the query component 31 and/or the table component 30, or by custom coding an event handler in script or a language which bridges the two components. However, the present invention enables this connection to be automatically created at run time, so as to eliminate the need for such custom coding. This feature, which is described below, enables Web designers with little or no programming experience to easily construct Web-based applications.

In accordance with the present invention, the components 26 also include a special-purpose component 29 referred to as the "director" component, as shown in Figure 4A. The director component 29 provides various control functions associated with the application, which are performed in conjunction with the infrastructure 28. These functions include maintaining registrations for the other components 26 and maintaining certain global default parameters of the application 19. The director component 29, the table component 30, and the query component 31 each have a connection to the infrastructure 28.

Figure 4B illustrates how a software application created according to the present invention may be distributed across multiple tiers (two in this case). As shown, the infrastructure on one tier may have a connection to a corresponding infrastructure on at least one other tier. Connections 47 and 49 between application components and connection 33 are established across tiers. A director component and an infrastructure are instantiated on each tier on which a component is to execute, and more specifically, for each different session on each tier on which a component is to execute. Thus, if an application includes both client-side components and server-side components, then separate director components and infrastructure are instantiated on the client and on the server and, more specifically, on any client or server machine on which a component is to execute. In the case of a multi-tiered application, the director component 29 is connected to a corresponding director component on at least one other tier via the infrastructures.

Referring still to Figure 4B, the application 19 is distributed between the client 20a and the application server 24a. An infrastructure 28a is instantiated on the client 20a, while an essentially identical infrastructure 28b is instantiated on the server 24a. In this example, the

components 26 of application 19 include a query component 31, a chart component 45, and a cube component 46, each instantiated on the client 20a and connected to the infrastructure 28a. The components 26 further include a director component 29b and a table component 30, both instantiated on the server 24a and connected to the infrastructure 28b. The query component 31 is for enabling a user to input a query upon a database (not shown). The table component 30 is for communicating with the database and for maintaining a result set based on the query. The cube component 46 is for transforming two-dimensional data from the database into a multi-dimensional representation which supports examining slices or cross-sections of the data. The chart component 45 is for displaying the results to a user. The client 20a and server 24a further maintain communications infrastructures 51a and 51b, respectively, which manage communications over connection 33. Communications infrastructure 51a has a connection to the client's main infrastructure 28a, while communications infrastructure 51b has a connection to the server's main infrastructure 28b and (via link 33) to communications infrastructure 51a.

In this example, connection 47 is automatically established between the query component 31 and the table component 30 at run time, to allow the query component 31 to specify to the table component 30 data that is to be retrieved from the database. Connection 48 is automatically established between the chart 45 and the cube 46 at run time, to allow the chart component 45 to receive data from the cube component 46. Further, connection 49 is automatically established between the table component 30 and the cube component 46 at run time, to allow data to be transferred from the table 30 to the cube 46.

The process for invoking an application will now be described at a high level with reference to Figure 5. In step 501, the user of client 20a (Figure 3) clicks on a hypertext link to one of the HTML pages 25. The client 20a then transmits a request for the corresponding page to the HTTP server 21a in step 502. In step 503, the HTTP server 21a passes the request on to the application server 24a, and in step 504, the application server 24a retrieves the requested page from the storage facility 27.

Each of the components 26 is referenced by an applet tag in at least one of the HTML pages 25. Each applet tag may include information that is not found in conventional applet tags, such as a designation of which tier the applet is to be executed on and/or the scope (duration) of a component's instantiation. Therefore, in step 505, the application server 24a determines whether the requested page refers to any server-side components (i.e., components to be executed on any tier other than the client tier) based on the contents of the corresponding applet tags and any rules or heuristics established by the system administrator. If no server-side components are identified in the requested page in step 505, then in step 506 the application server 24a passes the requested page to the HTTP server 21a, which passes the page to the client 20a in step 507. If the application server identifies one or more server-side components referenced by the requested page in step 505, then the application server 24a performs the server-side routine of step 508, which is described below. Following the server-side routine, the routine of Figure 5 proceeds from step 506.

As noted above, each applet tag may include information which is not found in conventional applet tags, such as a designation of which tier the component is to be executed on and/or the scope (duration) of a component's instantiation. This information is provided in the form of

applet attributes designated "tier" and "scope", respectively. Possible values for "tier" include, for example, "client", which is the default in one embodiment, "server", or other values to specify more complicated designations, as set up by the system administrator. In addition, a component can be specified for a particular client tier or a particular server tier by defining additional values for "tier", or a component can support a dynamic choice of tier by the application server 24a, as noted above. The ability to explicitly assign a component to a given tier allows the application designer to determine where to put the application logic and how to distribute the processing load. By using the tier attribute, the difference between running on the client or the server becomes a simple property that can be easily toggled for testing purposes or conditionally set to run on a client or server based on user profile, browser specifications, or other run-time properties, such as server load, performance, etc.

The "scope" attribute represents the duration of a component's instantiation on its specified tier and can have the values "request", "application", "session", or "global". "Request" scope, which is the default in one embodiment, indicates that the corresponding component will be instantiated as an object only while the corresponding HTML page is rendered. "Application" scope indicates that the corresponding component will be instantiated and available to other tiers provided they are executing the same application. "Session" scope indicates that the component will be instantiated and available only throughout a particular browser session, and "global" scope indicates that the applet is to remain instantiated until it is explicitly destroyed and is accessible to all clients and applications.

An example of an applet tag including the "tier" and "scope" parameters is as follows:

```
<applet code=sample tier=server scope=session  
name=sample_component>  
</applet>
```

In the above example, the component known as "sample_component" is specified to execute on a server tier and to have "session" scope.

Generally, when a page is requested, a number of objects are instantiated (created) for each component referenced by the page. Instantiation of a component is performed by the tier to which the component has been assigned. Thus, in the example of Figure 4, the client 20a instantiates the table component 30 and the query component 31 when the page referencing these components is downloaded. Also, downloading the first HTML page that references one of the components 26 to the client also results in the creation of the director component 29.

Upon instantiation of a component, the component registers its existence, as will now be explained with reference to Figure 6. Figure 6 illustrates a routine by which a component is instantiated and registered. In step 601, the client 20a receives a requested HTML page. In step 602, if the client 20a detects an applet tag for a component, the browser instantiates the component in step 603. In step 604, the infrastructure 28 registers that component's existence, and the browser renders the page in step 605. If no applet tag was found for a component in step 602, then the client 20a simply renders the page normally in step 605. When the first HTML page containing an applet tag for one of the components 26 is downloaded to the client 20a, the director component 29 and each other

component referenced by that page (e.g., table component 30 or query component 31) are instantiated on the client computer system 1.

Referring now to Figure 7, the components 26 and infrastructure 28 will now be described in greater detail. The infrastructure 28 includes an object referred to as the director sub-component 34. The director component 29 generally provides the front end to director sub-component 34, in addition to the functions mentioned above. The director sub-component 34 is responsible for managing the core functions of the application 19, including maintaining registrations of the other components (e.g., table component 30, query component 31, etc.) associated with the application 19 and handling requests from various components for connections to other components. In particular, (execution of) the director sub-component 34 maintains a list of registrations of all components of the application 19, which includes a reference to each component, and the use of which is described below. These functions of the director sub-component 34 enable the components to automatically acquire connections to each other without the need for the application developer to explicitly "wire" (i.e., write code to specify) the connections.

Note that in one embodiment, an application of the present invention will create one director component and one director sub-component for each distinct scope level used by the application. For example, if certain components default to the "request" scope while others are specified as "application" scope for a given tier, then at least two director components and at least two director sub-components would be created on that tier, i.e., one for "request" scope and one for "application" scope.

In addition to instantiating components 29, 30, and 31, the client 20a also instantiates several additional objects. (If one or more components are server-side components, however, then these additional objects may be created by the application server 24a.) The additional objects include, for each component, an object referred to as a context sub-component. Each context sub-component is associated with a particular one of the components 26, although the context sub-components are actually part of the infrastructure 28. Each context sub-component may be implemented as a Java class and includes methods for setting parameters and properties, locating other components, and registering its corresponding component. Hence, the client 20a (or application server 24a) creates a context sub-component 35a for the director component 29, a context sub-component 35b for the table component 30, and a context sub-component 35c for the query component 31. Each context sub-component has a reference to its corresponding component and a reference to the director sub-component 34. Note that, although the illustrated embodiment includes a separate context sub-component for each component, the context sub-components can be implemented as a single object which has a reference to all components.

Figure 8A illustrates a routine that is performed as a result of executing the objects shown in Figure 7. In step 801, the client 20a (or application server 24a in the case of the server-side routine--see below) finds an applet tag for one of the components 26 upon receipt of the HTML page and instantiates that component. Next, in step 802 the newly created component creates its context sub-component. For example, if the browser initially finds the tag for the table component 30 in step 801, then in step 802 the newly created table component 30 creates its context sub-component 35b.

Since the tag for the director component 29 may not be the first tag identified by the client 20a, the director component 29 may not be the first component to be instantiated. Consequently, the first component 26 to be instantiated during execution of the application 19 will then instantiate the director sub-component 34 (using its context sub-component). Thus, in step 803, the newly instantiated context sub-component attempts to locate the director sub-component 34. If the component is the first component to be instantiated, then its context sub-component will not be successful in locating the director sub-component 34 (step 804). In that case, the context sub-component first creates the director sub-component 34 in step 806 and then registers its newly instantiated component with the director sub-component 29b in step 805. If the director sub-component was found in step 804 (i.e., the component is not the first component to be instantiated), then the context sub-component registers its component with the director sub-component 34 in step 805. The above routine of Figure 8A is repeated for each component referenced in the downloaded HTML page.

As noted above, certain components may be specified as server-side components using the "tier" attribute of the applet tag. When that is the case, the application server 24a performs the server-side routine of step 508 (Figure 5). The server-side routine of step 508 involves a process substantially similar to that of Figure 8A, discussed above. That is, any server-side components and their supporting infrastructure (i.e., director sub-component and context sub-component) are instantiated on the appropriate server tier, which may be the application server 24a or any other server in the network. In addition, the application server 24a strips out the corresponding applet tag from the HTML page before allowing the page to be returned to the client 20a. As a result, the client 20a sees

only a standard HTML page. If the output of a server-side component is HTML code, then that code can be rendered back into the calling page by the application server 24a before the page is provided to the client 20a. Note that when a component is instantiated on a server tier and must persist for longer durations than handling the current request, its scope will generally be altered from the default "request" scope to a scope of longer duration.

Note that the present invention enables different components of an application, and even different components referenced by a single HTML page, to be instantiated and executed on different tiers. In addition, the present invention enables components of a single application to be concurrently instantiated and executed on three or more different tiers, e.g., a client tier and two or more server tiers. Figure 8B illustrates a simple example of these capabilities. Another example is described below with reference to Figure 10. In Figure 8B, an HTML page 62 and three executable components 63, 64, and 65 referenced by the HTML page 62 are parts of an application of the present invention and are stored in storage facility 27. Of course, components 63, 64, and 65 are not instantiated as stored in storage facility 27.

The application server 24a receives a request for the HTML page 62 from the client 20a via the HTTP server 21a (see Figure 3). In response to the request, the application server 24a provides the HTML page 62 to the client 20a. Also in response to this request, the application server 24a causes component 63 to be instantiated and executed on the client 20a, causes component 64 to be instantiated and executed on the application server 24a, and causes component 65 to be instantiated and executed on another server 66.

The tier assignments of components 63, 64 and 65 may be based on any or all of the above-mentioned tier assignment techniques, such as use of the tier attribute, load balancing, performance, etc. Thus, in this example, components 63b, 64b, and 65b are concurrently instantiated and executed on three separate tiers, the client 20a, the application server 24a, and the second server 66. Note, however, that components need not be executed on the application server 24a; that is, components could just as easily be executed on the client tier 20a and two or more server tiers other than the application server 24a.

Many application components may be considered to be data "producers", data "consumers", or "data controllers". A functional connection is often required between a particular data producer and a particular data consumer or between a data producer and a data controller, to allow data to be communicated between those components. As noted above, the director sub-component 34 (Figure 7) maintains a list of registrations of all application components 26, which is used to automatically establish such connections. In accordance with the present invention, connections between executable components are established automatically and asynchronously, as required by the application 19, in response to requests from data producers or data consumers. Consequently, individual components effectively "snap together" automatically, without an application developer having to specify the connections.

Referring again to the example of Figure 7, a connection 32 is required between the table component 30, which is a data producer, and the query component 31, which is a data controller. The function of establishing connections between components is generally performed by the infrastructure 28 and, more specifically, by the context sub-component

of whichever component is requesting a connection. The process of establishing a connection between two components is described now in greater detail with reference to Figure 9.

In step 901, a component, such as a data consumer, provides its context sub-component with a request. The request specifies a set of criteria to be used to identify one or more appropriate "service providers", or components to which the requesting component can be connected, such as a data producer. The criteria may include, for example, the name of one or more components to which a connection is needed or the name of one or more interfaces which the requesting component supports. In this context, an "interface" is a specification of the methods, or subroutines, which a component supports.

The request is actually provided to a method of the context sub-component Java class designated "findServiceProvider." (See the sample Java code below for an example of how the method findServiceProvider is used.) Thus, in step 902, the method findServiceProvider contacts the director sub-component 34 and requests the list of registered components maintained by the director sub-component 34. In step 903, the director sub-component 34 provides the list of registrations to the findServiceProvider of the requesting context sub-component. In step 904, findServiceProvider searches the list for the appropriate service provider or providers. Specifically, findServiceProvider uses the references in the list to identify any component or components which match the criteria specified in the request (e.g., which have the specified name or support the specified interface). Note that Java allows an object to be queried to determine whether the object supports a particular interface.

More than one of the other components may satisfy a request. Accordingly, the criteria specified in the request also include a flag, which the requesting component can set to indicate that it requires a reference to all components which satisfy its request. In addition, the criteria of the request may specifically limit the request to service providers (components) on the same tier as the requesting component (the local tier) or to service providers on a remote tier. Furthermore, remote requests can be limited to a set of scopes, etc.

Thus, if `findServiceProvider` finds a component or multiple components which satisfy the criteria in the request (step 905), then in step 906 `findServiceProvider` provides to the requesting component the reference or references to the one or more components which satisfy the request. If no component is found which satisfies the request, then the routine ends, with the request remaining pending until a component which satisfies the request becomes instantiated and registered on the current tier. Thus, connections between the components 26 of the present invention are established automatically and asynchronously as required by the application 19. "Asynchronously" means, in this context, that the requesting entity is able to perform other processing functions while the request is pending, rather than having to wait for a response to the request. Consequently, there is no specific temporal relationship between the issuance of a request for a connection and the response to that request; a request can be satisfied at any time after it has been submitted. Further, a connection can be automatically established even if the component which satisfies the request is not yet instantiated at the time the request is submitted.

In addition, a component may become "unregistered" for various reasons. Unregistration of a component may occur, for example, when

the component becomes no longer instantiated, such as if the component is explicitly destroyed. If a component becomes unregistered for any reason, then the director sub-component 29a provides any components which have a connection to the unregistered component with a notice of unregistration, which indicates to such components that the unregistered component is no longer available.

Below is an example of the Java code for a HTML page that incorporates two application components, designated Blox1 and Blox2, as well as the Java code for these components. The component Blox1 supports an interface, ITest, which contains one method, getResponse. The method getResponse returns a string representation of the text in the text box displayed in this component. The function of these two components is to enable a user to change the text in the text box, such that another component can retrieve the text using the ITest interface. The component Blox2 has a user interface which shows certain status information in a text box and supports a button which, when pushed, will invoke the getResponse method of the ITest interface on Blox1. Blox2 locates this interface using the method IBloxContext.findServiceProvider. The call to this method may be in the bloxStart method of the component, which is discussed below. When the button is pushed, the string value, which originated in the text box of Blox1, is displayed in the text box of Blox2.

Thus, the sample HTML page is as follows:

```
<html>
```

```
<head>
```

```
<title>Example application</title>
```

</head>

<body>

<h1>Example application</h1>

<p>

Here are three Components: (Director, Blox1, and Blox2)

Director and Blox2 will be colored magenta because the Application.BGColor property set on Director will be used as the default BGColor property for all components that don't have their own BGColor property. Note that Blox2 doesn't define any BGColor, and gets its default value from Director.

Blox1 will be colored cyan since it overrides the default BGColor property.

When this page finishes loading, "found ITest interface" will be displayed in the text field in Blox2. This means that the findServiceProvider call in Blox2 located the ITest interface defined in Blox1.

If one enters text in the text box in Blox1 and then push the button in Blox2, then the text that you entered in Blox1 will be displayed in Blox 2.

Note that no explicit wiring up between Blox1 and Blox2 occurs in the HTML file.


```
<applet code=alphablox.blox.applet.Director
        codebase="/classes"
        name=Director
        height=20
        width=20>
```

```
<param name=Application.BGCOLOR value=magenta>
```

```
</applet>
```

```
<applet code=Blox1
        codebase="/classes"
        name=Blox1
        height=200
        width=200>
```

```
<param name=BGCOLOR value=cyan>
```

```
</applet>
```

```
<applet code=Blox2
        codebase="/classes"
        name=Blox2
        height=200
        width=200>
```

```
</applet>
```

```
</body>
```

```
</html>
```

The following is sample Java code for the component Blox1.

```
import alphablox.mortar.*;
import alphablox.awt.*;
import java.awt.*;

public class Blox1 extends BloxApplet implements
ITest
{
    TextField textField = new TextField("default value"); // a text box where
one can edit the response

    public void bloxInit()
    {
        super.bloxInit();    // make sure
BloxApplet.bloxInit gets an opportunity to run

        setLayout(new BorderLayout());
        add("Center", textField);
    }
}
```


// bloxStart was not overridden, but the base class BloxApplet defined a bloxStart that caused this component to become registered with the infrastructure.

```
// implementation of ITest interface

/**
 * Return the contents of the text field
 */
public String getResponse()
{
    return textField.getText();
}
}
```

The following is the sample Java code for Blox2. The sample code illustrates the use of methods bloxStart and findServiceProvider.

```
import alphablox.mortar.*;
import alphablox.awt.*;
import java.awt.*;

public class Blox2 extends BloxApplet implements
IServiceFoundListener
{
    TextField textField = new TextField(20); // a text box where one
can write text
    Button pushMe = new Button("PushMe"); // a button which one can
push to cause an action to occur
```

```
private ITest test;    // where the ITest interface will be put
when it is received from Bloxl

public void bloxInit()
{
    super.bloxInit();    // make sure BloxApplet.bloxInit gets an
opportunity to run

    setLayout(new BorderLayout());
    add("Center", textField);
    add("South", pushMe);

    pushMe.disable();    // disable this button until service
provider is found
}

public void bloxStart()
{
    super.bloxStart();    // make sure BloxApplet.bloxStart gets an
opportunity to run

    getBloxContext().findServiceProvider("ITest", this); // start search
for any component that supports the ITest interface
}

/**
```

```
* This method is invoked by infrastructure when a service requested by  
findServiceProvider is satisfied  
*/  
public void serviceProviderFound(ServiceEvent event)  
{  
    textField.setText("got a response from findServiceProvider");  
  
    Object serviceProvider = event.getServiceProvider();  
    if (serviceProvider instanceof ITest) // it was indeed an ITest supporting  
object  
    {  
  
        test = (ITest)serviceProvider; // therefore, remember this interface . .  
  
        textField.setText("found ITest interface");  
  
        pushMe.enable(); // and enable the button so it  
can be called  
    }  
}  
  
/**  
  
* This method is invoked by the infrastructure when a service provider  
is no longer available.  
* We ignore this event.  
*/
```

```
public void serviceProviderUnregistered(ServiceEvent event)
{
}

/**
 * This method is called by AWT (windowing library) when, for example,
 a button is pushed.
 */
public boolean action(Event event, Object arg)
{
    boolean handled = false;

    if (event.target == pushMe) // when the pushMe button is pressed
    {
        String response = test.getResponse();           // ask Bloxl for a
"response"
        textField.setText("got \" + response + "\" from Bloxl"); // display
what was returned
        handled = true;                               // tell AWT that
event was handled
    }
    return handled;
}
}
```

The following is sample Java code for the ITest interface:

```
public interface ITest
```

```
{  
    String getResponse();  
}
```

As is well known, Java applets support four standard methods (among others): `init`, `start`, `stop`, and `destroy`. In the present invention, these methods are called by the client 20a or the application server 24a to manage the lifetime of an applet (e.g., one of the components 26). Generally, calling the `init` method causes an applet to be initialized; calling `start` causes an applet to be started; calling `stop` causes an applet to be stopped; and calling `destroy` causes an applet to be destroyed. In the present invention, these four methods call four additional methods, referred to as `bloxInit`, `bloxStart`, `bloxStop`, and `bloxDestroy`. For a given component (e.g., `Blox1` or `Blox2` in the example below), the default (base) implementation of `bloxInit` causes the component's context sub-component to be created. The default implementation of `bloxStart` causes the component to be registered with the director sub-component. In addition, if there are any requests for connections pending from other components, the act of registration causes a determination of whether the newly registered component matches the criteria of any of the requests and, if so, a message to be sent to the service requestor that the component is registered. The default implementation of `bloxDestroy` method causes a component to be unregistered and a message of unregistration sent to any component which has a connection to that component.

In the example above, `Director` is the embodiment of the director component 29. `Director` serves only as a place to define certain global defaults, as noted above. An example of such a default is the definition of

Application.BGCOLOR property on Director, which is used by all components of the application as the "background color" unless explicitly overridden on individual components.

Blox1 and Blox2 are both derived from the base class BloxApplet. This class is defined in the infrastructure 28 (see Figures 3, 4A, and 4B). The class BloxApplet defines the default implementation of bloxInit, bloxStart, bloxStop, and bloxDestroy, described previously. BloxApplet creates the BloxContext, which in turn exposes the IBloxContext interface. The getBloxContext method in BloxApplet returns the IBloxContext interface.

Blox1 and Blox2 override the default bloxInit implementation, while delegating to it to get the base functionality. In the overridden method, each blox does its specific initialization. Blox1 does not override bloxStart, bloxStop, and bloxDestroy, and therefore simply gets the default implementation from BloxApplet. Blox2 overrides the default bloxStart implementation, while delegating to it to get the base functionality. In the overridden method, Blox2 attempts to locate an object which supports the ITest interface. It does so by using the getBloxContext method in BloxApplet to get the IBloxContext interface, and then calls the findServiceProvider method on this interface.

There are two possibilities in this example: Either Blox1 will be instantiated first or it will be instantiated second. If Blox1 is instantiated first, then when its base implementation of bloxStart is called, it will be registered with the Director sub-component 34. When Blox2's bloxStart is called, the findServiceProvider call will result in the infrastructure's locating Blox1 and invoking Blox2's serviceProviderFound method, notifying Blox2 of the component which satisfies its request.

If Blox2 is instantiated first, the findServiceProvider call will not find any component to satisfy the request, and the request will remain pending. When Blox1 is instantiated, and when its base implementation of bloxStart is called, it will be registered with the Director sub-component 34. This action will result in the infrastructure's noticing the pending request and notifying Blox2 (via the serviceProviderFound method) of Blox1's existence.

As mentioned above, the present invention also provides the capability to distribute components of a single application across three or more separate tiers (e.g., a client tier and two or more server tiers) and to automatically establish connections between the appropriate components. An example of this capability was described above with reference to Figure 8B. The following is sample Java code for another example in which components are executed on three separate tiers, which is also illustrated in Figure 10.

```
<html>
<head>
<title>N-tier Example</title>
<body>
<h1>N-tier Example</h1>

<body>
<applet code=alphablox.blox.applet.Director
      name=Director
      codebase=/classes
      tier=client
```

```
<param name=ApplicationName value=MyApp>
```

```
</applet
```

```
<applet code=alphablox.blox.applet.Table
```

```
  name=Table
```

```
  codebase=/classes
```

```
  tier="server:host=Host1"
```

```
  scope=application
```

```
<param name=ApplicationName value=MyApp>
```

```
</applet>
```

```
<applet code=alphablox.blox.applet.Cube
```

```
  name=Cube
```

```
  codebase=/classes
```

```
  tier="server:load<100,class=olap"
```

```
  scope=application
```

```
<param name=ApplicationName value=MyApp>
```

```
</applet>
```

```
<applet code=alphablox.blox.applet.Grid
```

```
  name=Grid
```

```
  codebase=/classes
```



```
        tier=client

<param name=ApplicationName value=MyApp>

</applet>

</body>
</html>
```

In the above "n-tier" example, there are four boxes involved: Director, Table, Cube, and Grid, as shown in Figure 10. Director and Grid are specified to be run on the client tier. Table and Cube are specified to run on server tiers. Note that while Director is explicitly authored to execute on the client tier, in accordance with the present invention, additional Director components are also automatically instantiated and executed on each additional tier on which another component is to execute. Table is explicitly specified to run on a server named "Host1". Cube is specified to run on any server tier which match specific criteria. These criteria are configurable by system administrators and are specified here only for purposes of illustration. In this example, Cube will be run on any available server whose "load" factor is less than 100 and which belongs to a class of servers that have been labeled as "olap". Both of the server components are specified as being of "application" scope, which means that these components will be available to any other components in the same application. The application used for all components in this example is named, "MyApp".

In operation, the application server 24a (not shown in Figure 10--see Figure 3) will strip out the two server-scope components from the

HTML sent to the client 20a. Thus, the client 20a will only instantiate the Director and Grid components. Prior to serving the page to the client 20a, however, the application server 24a will locate the appropriate servers which match the specified tier criteria and cause these components to be executed on those tiers. It will also keep track of all the servers involved in serving this request so that it can know which servers to contact to satisfy findServiceProvider requests.

The Table component, once instantiated on its tier, will locate its external data source (due to the role that Table plays, it doesn't attempt to locate another component, but uses external protocols to determine its data source). The Cube component, however, cannot function until it locates a Table datasource (more accurately, a component which supports the ITableProducer interface, which Table does). Cube attempts to locate the ITableProducer component on its own tier, and doesn't find it. It then contacts the application server 24a to find the correct component. Because the application server 24a knows that it instantiated components on Host1 and the requesting tier, it passes the request to Host1 (assuming that the requesting tier has already been searched). Host1's Director gets the request and locates the Table component, returning this result to the application server 24a. The application server 24a then returns this result to Cube. A connection between Cube and Table is thus established.

On the client tier, Grid looks for its data source (of type ICubeProducer). Not finding the data source on its tier (client 20a), it passes the request to the application server 24a. The application server 24a passes the request to both Host1 and the host on which it started the Cube component. The Cube component (which supports the ICubeProducer interface) is thus located, and a reference to it returned to the Grid component on the client 20a.

Implicitly, remote proxies and stubs are created to allow the appearance of a local interface when, in fact, the interface spans separate tiers.

The present invention also includes a mechanism for generating applications of the type described above, based on templates. Using this mechanism, any conventional Web page design environment or Internet integrated development environment (IDE) (e.g., Microsoft FrontPage, Netscape Visual JavaScript, or Net Objects Fusion) can be used to create applications according to the present invention. Specifically, a conventional storage facility can be used to store a library of components that can be included by application designers in various different Web applications. Thus, an application designer can select particular components to create custom-designed applications. This process may be performed via a graphical user interface such as illustrated in Figure 11. Figure 11 shows a display 35 including two windows 36 and 37. A number of icons 40b are displayed in window 37, each of which represents a component according to the present invention. The application developer uses a pointing device to place a cursor 38 on particular ones of icons 40b and to drag-and-drop them into window 36 to incorporate the corresponding components (i.e., Java classes) into the application. The interface of Figure 11 may be such that icon 40a representing the director component 29 appears in window 36 by default, since each application includes at least one director component.

In response to the application developer's dragging and dropping an icon into window 36, a corresponding applet tag is automatically inserted into the HTML page being authored. In one embodiment, window 36 can also be used to input text using a conventional HTML

editor, such that the application developer can graphically combine standard HTML with various components. In response to icons 40b being dropped within window 36, the corresponding components are effectively included within the application that is being created. Consequently, at run time the infrastructure 28 will automatically ensure that the correct connections between the components are automatically established. The application developer is not required to specify these connections.

Figure 12 illustrates another graphical user interface by which an end user can both assemble and execute a customized application. In this example, the application to be assembled is a database application. The display screen 50 is an HTML page, and the application includes one or more executable components according to the present invention. From the display screen 50, the end user can effectively select which components are included in the application and thereby generate a customized output report of data stored in the database. The display 50 also includes a set of options from which the user can select to generate a customized output report. The options include the particular format of the report to be generated. The user's selections from among these options determine which executable components are included in the application for the current run of the application.

In this example, the executable components which may be included in the application (which are not shown) include a grid component for generating a report in grid format, a chart component for generating a report in chart format, and a cube component for maintaining a multi-dimensional data set. Each of rectangles 55 represents an image depicting a reduced-size example of the corresponding format. Accordingly, the user can select one of buttons 54A to request a specific grid format;

alternatively, the user can select one of buttons 54B to request a specific chart format. In this example, there are two different grid formats and two different chart formats from which the user may choose. When the user subsequently clicks on the "Assemble" button 59, the user's format selections cause the appropriate component or components (i.e., grid component or chart component) to be automatically identified and executed. The user's selection of button 62 for "Multi-Dimensional Analysis" will result in the cube component being automatically included in the application and executed when "Assemble" button 59 is clicked on. Further, by employing the above-described techniques for automatically establishing connections between components, the appropriate connections are automatically established between the selected components in response to the "Assemble" button 59 being clicked on. The end user does not need to wire any of the connections.

The end user can also customize various attributes of the application components using screen 50. For example, the user can: enter the name of the application to be assembled in field 51 or choose from a list of previously assembled applications; specify a predefined data set in field 52 rather than specifying a new query; specify text to be included in the header, body, and footer of the report by entering the text in fields, 53, 56, and 57, respectively; and, specify a background color for the report in field 58.

Thus, using a graphical user interface such as that shown in Figure 12 in conjunction with other features of the present invention, the user can assemble and execute a customized application by selecting only the desired components for inclusion in the application. Connections between the selected components are automatically established when the

application is executed, such that the end user does not need to be concerned with "wiring" components together.

Thus, a tier-neutral development mechanism for Web-based applications has been described. Although the present invention has been described with reference to specific exemplary embodiments, it will be evident that various modifications and changes may be made to these embodiments without departing from the broader spirit and scope of the invention as set forth in the claims. Accordingly, the specification and drawings are to be regarded in an illustrative rather than a restrictive sense.

CLAIMS

What is claimed is:

1. A mechanism for creating a software application for execution on a network having a plurality of tiers, the mechanism comprising:
 - specifying means for specifying application components, such that any application component specified using said specifying means can be assigned to execute on any of the plurality of tiers of the network; and
 - means for enabling the application components to be associated with at least one hypertext-based page, such that the components are executable in response to requests for the at least one hypertext-based page.
2. A mechanism according to claim 1, wherein the specifying means comprises means for specifying a plurality application components of a single software application for execution on at least three separate tiers of the network.
3. A mechanism according to claim 1, wherein the plurality of application components of the single software application comprises a first component specified to execute on a first tier and a second component specified to execute on a second tier different from the first tier, and wherein the mechanism further comprises means for enabling communication between the first component executing on the first tier and the second component executing on the second tier without a user having specified connections between the components.

4. A mechanism according to claim 1, wherein the specifying means comprises means for specifying a tier on which a first application component is to execute using an attribute of a tag referencing the first application component.
5. A mechanism according to claim 1, further comprising means for dynamically determining, based on a current state, a tier on which a first application component is to execute in response to a request for a hypertext-based page corresponding to the first application component.
6. A mechanism according to claim 1, wherein each of the application components is an object.
7. A method of executing an application on a network having a plurality of tiers, the application including at least one non-executable component and a plurality of executable components referenced by said at least one non-executable component, the plurality of executable components including a first executable component, a second executable component, and a third executable component, the plurality of tiers including a first tier, a second tier, and a third tier, the method comprising the steps of:
 - receiving from a client system at least one request for said at least one non-executable component;
 - in response to said at least one request:
 - determining on which tier of the network each of the first, second, and third executable components should be executed;
 - causing the first executable component to be executed on the first tier;

causing the second executable component to be executed on the second tier; and

causing the third executable component to be executed on the third tier.

8. A method according to claim 7, wherein each of the at least one non-executable components is a hypertext-based page.

9. A method according to claim 7, wherein the first tier is the client system, the second tier is a first server system, and the third tier is a second system.

10. A method according to claim 7, wherein each of the first, second, and third executable components is an object.

11. A method according to claim 7, wherein the determining step comprises, for at least one of the first, second, and third executable components, determining on which tier of the network the component should be executed based on an attribute of a tag in a non-executable component.

12. A method according to claim 7, wherein the determining step comprises, for at least one of the first, second, and third executable components, dynamically determining on which tier of the network the component should be executed based on a current state.

13. A method according to claim 7, wherein each application component is defined according to a single mechanism by which an application component can be specified for execution on any tier on the network.

14. A method of enabling a target processing system to execute an application, the method comprising the step of transmitting sequences of instructions from a host processing system to the target processing system, the instructions for causing the target processing system to perform the method of claim 7 when executed by the target processing system.

15. A method creating a software application for execution on a network having at least three tiers, the method comprising the steps of:

specifying a plurality of application components of a single application, the plurality of application components including at least three application components, such that each of the at least three application components can be executed on a separate tier of the network;
and

associating each of the plurality of components with at least one hypertext-based page, such that each component is executable in response to a request for an associated hypertext-based page.

16. A method according to claim 15, further comprising the step of enabling communication between a first component executing on a first tier and a second component executing on a second tier without using a specification of a connection between the first and second components.

17. A method according to claim 15, further comprising the step of specifying a tier on which one of the plurality of application components is to execute.

18. A method according to claim 15, wherein the step of specifying a tier on which one of the plurality of application components is to execute comprises specifying a tier in an attribute of a tag in the hypertext-based page referencing said one of the plurality of application components.

19. A method according to claim 15, wherein each of the application components comprises an object.

20. A mechanism for creating an application for execution on a network having a plurality of tiers, the mechanism comprising:

means for specifying application components, such that application components of a single application can be concurrently executed on at least three separate tiers of the network; and

means for enabling the application components to be associated with at least one hypertext-based page, such that the application components are executable in response to at least one request for the at least one hypertext-based page.

21. A mechanism according to claim 20, wherein the application components of the single application comprises a first component for execution on a first tier, a second component for execution on a second tier, and a third component for execution on a third tier.

22. A mechanism according to claim 20, wherein the means for specifying comprises means for specifying the tier on which each application component is to execute.

23. A mechanism according to claim 20, wherein each of the application components is an object.

24. A method of creating a software application for execution on a network, the network having a plurality of tiers, the method comprising the steps of:

creating a set of hypertext-based pages, the set including at least one hypertext-based page, each page of the set for receipt and display by at least one of the plurality of tiers on the network; and

including a reference to at least one of a set of executable components within one of the set of hypertext-based pages, the set of executable components including at least one executable component, the reference including a tier indication specifying a tier on which the at least one executable component is to execute.

25. A method according to claim 24, wherein the tier indication may specify any tier of a network having at least three tiers.

26. A method according to claim 24, wherein said tier is a first tier of a network, wherein the set of executable components comprise a first executable component and a second executable component, the method further comprising the steps of:

including a second reference to the second executable component within said one of the set of hypertext-based pages, the second reference

including a tier indication specifying a second tier on which the second executable component is to execute; and

including a third reference to the third executable component within said one of the set of hypertext-based pages, the third reference including a tier indication specifying a third tier on which the third executable component is to execute.

27. A method of enabling a target computer to create a software application, the method comprising the step of transmitting sequences of instructions from a host computer to the target computer, the sequences of instructions including instructions which, when executed on the target computer, cause the target computer to perform the method recited in claim 24.

28. A method of executing a software application, the software application for execution on at least one tier of a network having a plurality of tiers, the software application including at least one hypertext-based page referencing at least one executable component, the method comprising the steps of:

receiving from a remote tier a request for a hypertext-based page referencing an executable component;

determining on which tier of the network the executable component should be executed in response to the request;

causing said executable component to be executed on an appropriate tier of the network in response to the request; and

providing the hypertext-based page to the remote tier in response to the request.

29. A method according to claim 28, wherein the determining step comprises determining on which tier of the network said component should be executed based on an attribute of a tag in the hypertext-based page referencing said executable component.
30. A method according to claim 28, wherein the determining step comprises dynamically determining, in response to the request, on which tier of the network said component should be executed, based on a current parameter associated with the network.
31. A method according to claim 28, wherein each application component is defined according to a single mechanism by which an application component can be specified for execution on any of the plurality of tiers of the network.
32. A method according to claim 28, wherein the step of causing said component to be executed on an appropriate tier comprises the steps of:
determining whether said component should be executed on the remote tier; and
causing the component to be executed before providing said hypertext-based page to the remote tier if said component should not be executed on the remote tier.
33. A method according to claim 28, further comprising the step of, if said component should not be executed on the remote tier, causing a result of executing said component to be incorporated into said hypertext-based page before providing said hypertext-based page to the remote tier.

34. A method according to claim 28, wherein the application further comprises at least one hypertext-based page referencing a second executable component and at least one hypertext-based page referencing a third executable component;

wherein the step of causing said executable component to be executed on an appropriate tier of the network comprises causing said executable component to be executed on a first tier of the network;

the method further comprising the steps of:

causing the second executable component to be executed on a second tier of the network; and

causing the third executable component to be executed on a third tier of the network.

35. A method of enabling a target processing system to execute an application, the method comprising the step of transmitting sequences of instructions from a host processing system to the target processing system, the instructions for causing the target processing system to perform the method of claim 28 when executed by the target processing system.

36. A mechanism for enabling execution of a software application, the software application for execution on at least one tier of a network having a plurality of tiers, the software application including at least one hypertext-based page referencing an executable component, the mechanism comprising:

means for receiving from a remote tier a request for the hypertext-based page referencing the executable component;

means for determining on which tier of the network the executable component should be executed in response to the request; and

means for causing said executable component to be executed on an appropriate tier of the network in response to the request.

37. A mechanism according to claim 36, further comprising means for providing the hypertext-based page to the remote tier in response to the request.

38. A mechanism according to claim 36, wherein the means for determining comprises means for determining on which tier of the network said component should be executed based on an attribute associated with said component.

39. A mechanism according to claim 36, wherein the attribute is included in a reference to said executable component, and wherein the reference is included within the requested hypertext page.

40. A mechanism according to claim 36, wherein the means for determining comprises means for determining on which tier of the network said component should be executed based on a state of the network.

41. A mechanism according to claim 36, wherein each application component is defined according to a single model by which an application component can be specified for execution on any tier on the network.

42. A mechanism according to claim 36, wherein the means for causing said component to be executed on an appropriate tier comprises:

means for determining whether said component should be executed on the remote tier; and

means for causing the component to be executed before providing said hypertext-based page to the remote tier if said component should not be executed on the remote tier.

43. A mechanism according to claim 36, further comprising means for causing a result of executing said component to be incorporated into said hypertext-based page before providing said hypertext-based page to the remote tier, if said component should not be executed on the remote tier.

44. A mechanism according to claim 36, wherein the application further comprises at least one hypertext-based page referencing a second executable component and at least one hypertext-based page referencing a third executable component;

wherein the means for causing said executable component to be executed on an appropriate tier of the network comprises means for causing said executable component to be executed on a first tier of the network;

the mechanism further comprising:

means for receiving a request for a hypertext-based page referencing the second executable component;

means for receiving a request for a hypertext-based page referencing the third executable component;

means for determining on which tier of the network each of the second and third executable components should be executed;

means for causing the second executable component to be executed on a second tier of the network; and

means for causing the third executable component to be executed on a third tier of the network.

45. A machine-readable program storage medium tangibly embodying sequences of instructions, the sequences of instructions for execution by a processing system to perform steps for enabling execution of an application on at least one tier of a network having a plurality of tiers, the application including at least one hypertext-based page referencing an executable component, the steps comprising:

receiving from a remote tier a request for the hypertext-based page referencing the executable component;

determining on which tier of the network the executable component should be executed in response to the request; and

causing said executable component to be executed on an appropriate tier of the network in response to the request.

46. A machine-readable program storage medium according to claim 45, the steps further comprising providing the hypertext-based page to the remote tier in response to the request.

47. A machine-readable program storage medium according to claim 45, wherein the step of determining comprises the step of determining on which tier of the network said component should be executed based on an attribute of said component, wherein the attribute is included in a reference to said executable component, and wherein the reference is included within the requested hypertext page.

48. A machine-readable program storage medium according to claim 45, wherein the step of determining comprises the step of dynamically determining on which tier of the network said component should be executed based on a current parameter associated with at least one tier of the network.

49. A machine-readable program storage medium according to claim 45, wherein each application component is defined according to a single model by which an application component can be specified for execution on any tier on the network.

50. A machine-readable program storage medium according to claim 45, wherein the step of causing said component to be executed on an appropriate tier comprises the steps of:

determining whether said component should be executed on the remote tier; and

causing the component to be executed before providing said hypertext-based page to the remote tier if said component should not be executed on the remote tier.

51. A machine-readable program storage medium according to claim 45, further comprising the step of causing a result of executing said component to be incorporated into said hypertext-based page before providing said hypertext-based page to the remote tier, if said component should not be executed on the remote tier.

52. A machine-readable program storage medium according to claim 45, wherein the application further comprises at least one hypertext-based page referencing a second executable component and at least one hypertext-based page referencing a third executable component;

wherein the step of causing said executable component to be executed on an appropriate tier of the network comprises causing said executable component to be executed on a first tier of the network; and

the steps further comprising:

causing the second executable component to be executed on a second tier of the network; and

causing the third executable component to be executed on a third tier of the network.

53. A system for enabling a software application to be executed on at least one of a plurality of tiers on a network, the system comprising:

a processor;

a first storage unit accessible to the processor and having stored therein a plurality of executable components, and

a second storage unit accessible to the processor and having stored therein a plurality of pages functionally connected by hypertext links, the plurality of pages incorporating the plurality of executable components;

wherein the processor is configured to:

receive a request for one of the pages from a remote tier on the network; and

access one of the pages and a corresponding first one of the executable components in response to the request;

identify an appropriate tier on which said first one of the executable components should be executed in response to the request; and

cause said first one of the executable components to be executed on said appropriate tier in response to the request.

54. A system according to claim 53, wherein the processor is configured to identify said appropriate tier based on a reference to said first one of the executable components in the requested page, the reference including an attribute specifying a tier on which said first one of the executable components should be executed.

55. A system according to claim 53, wherein the processor is configured to dynamically identify said appropriate tier based on a current parameter.

56. A system according to claim 53, wherein the processor is further configured to:

determine whether the first one of the executable components should be executed on the remote tier; and

cause the first one of the executable components to be executed before the requested page is provided to the remote tier, if it is determined that the first one of the executable components should not be executed on the remote tier.

57. A system according to claim 53, wherein the processor is further configured to cause the requested page to be provided to the remote tier after the first one of the executable components is executed if the remote tier is not identified as said appropriate tier.

58. A system according to claim 53, wherein execution of said component results in an output, and wherein the processor is further configured to incorporate the output into the requested page before the requested page is provided to the remote tier if said first one of the executable components should not be executed on the remote tier.

59. A system according to claim 53, wherein the application further comprises at least one hypertext-based page referencing a second one of the executable components and at least one hypertext-based page referencing a third one of the executable components;

and wherein the processor is further configured to:

receive a request for a hypertext-based page referencing the second one of the executable components;

receive a request for a hypertext-based page referencing the third one of the executable components;

determine on which tier of the network each of the second and third ones of the executable components should be executed;

causing the second one of the executable components to be executed on a second tier of the network; and

causing the third one of the executable components to be executed on a third tier of the network.

1/14




FIG. 1

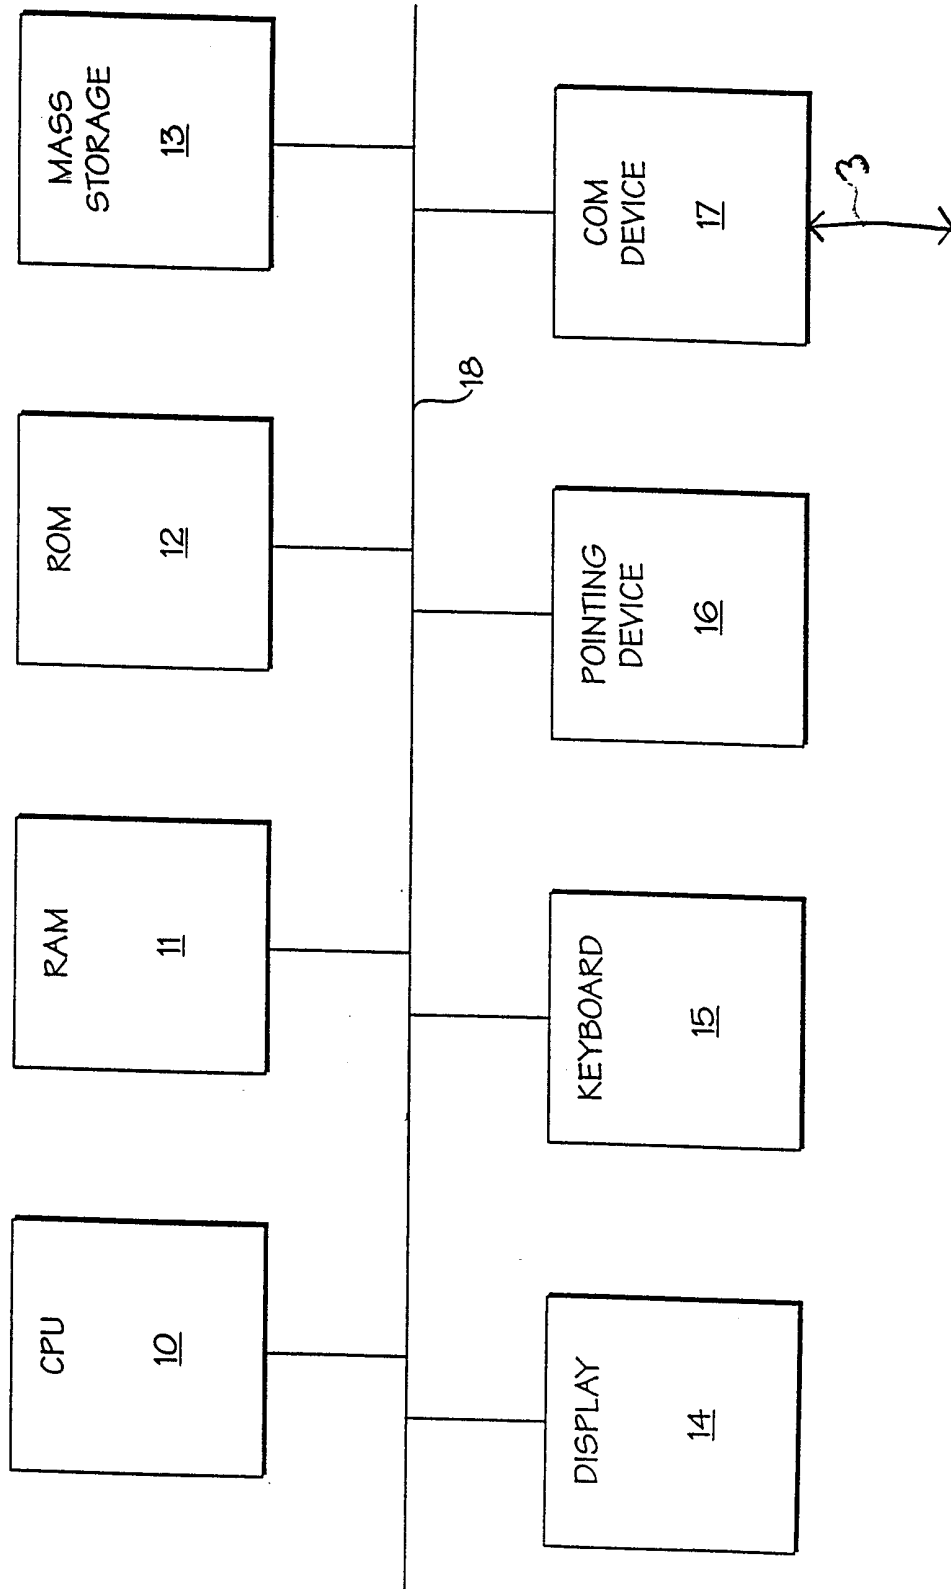


FIG. 2

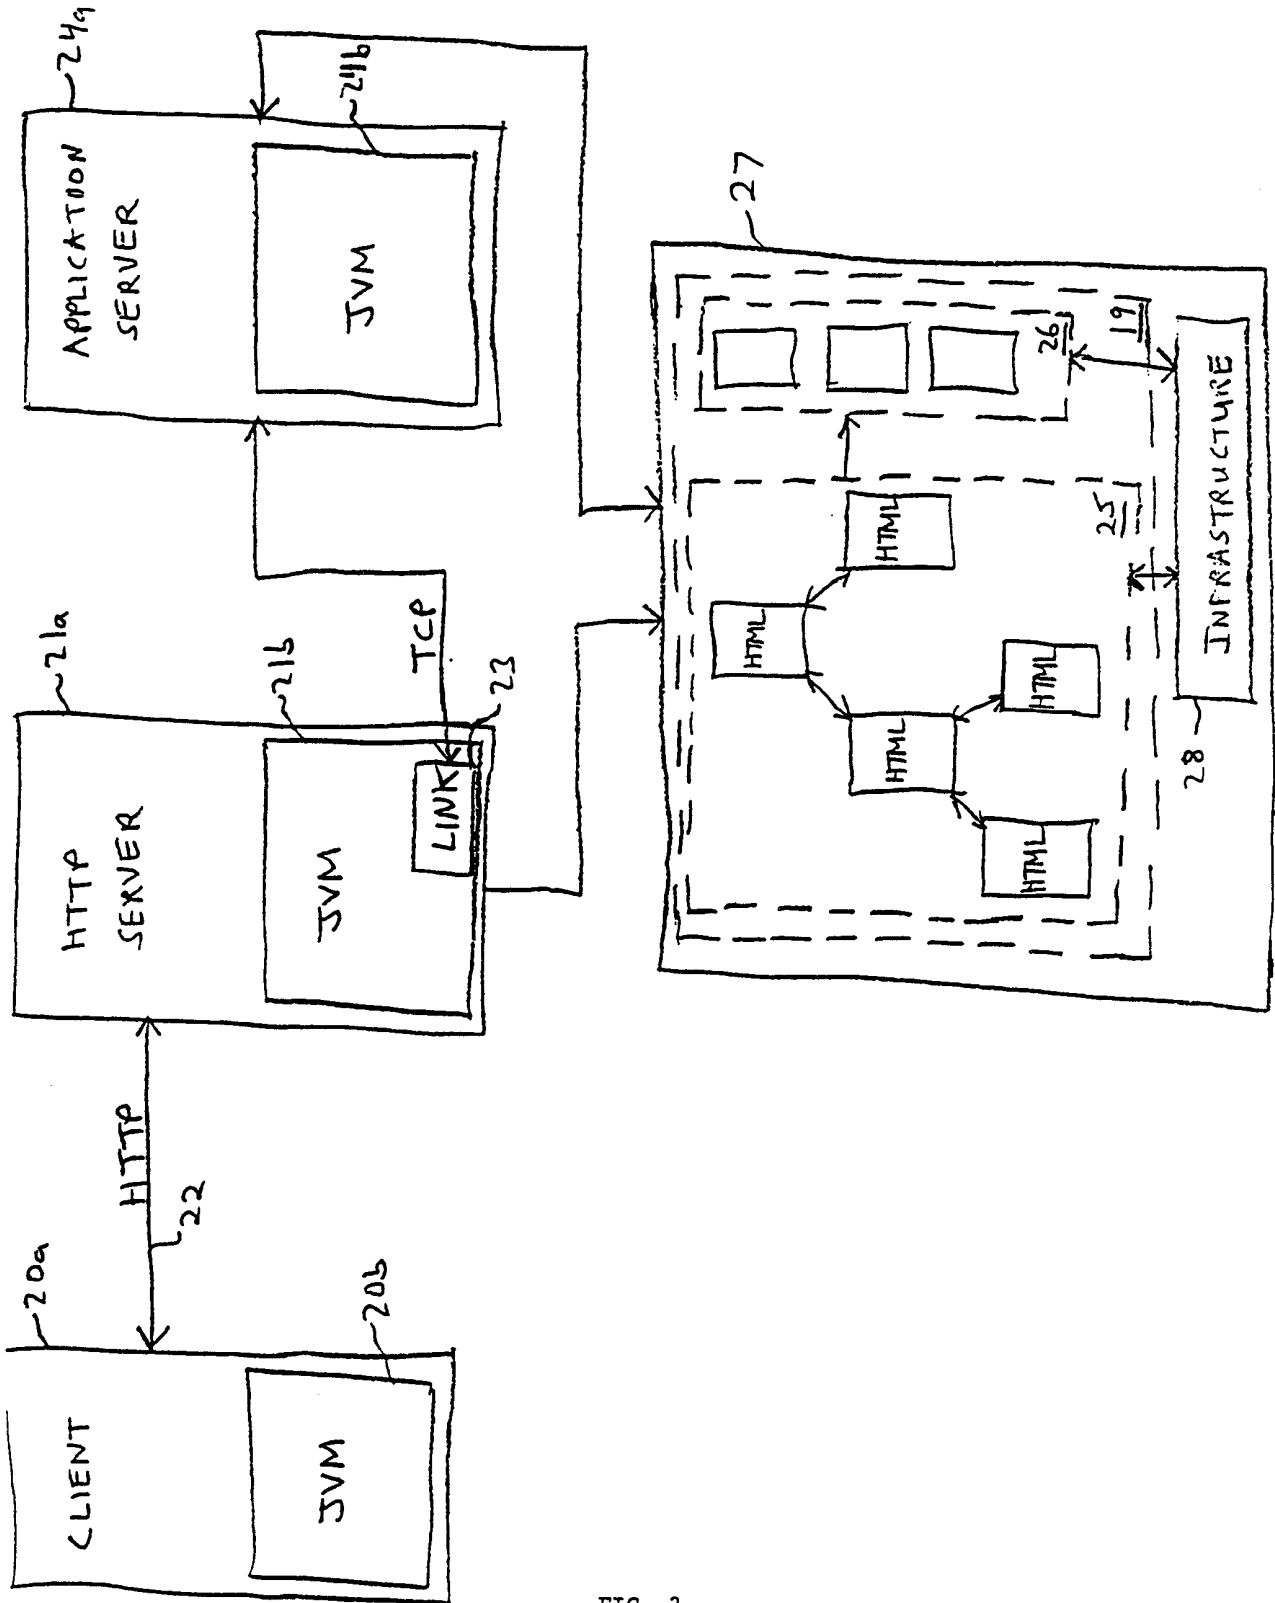


FIG. 3





FIG. 4A

FIG. 4B



6/14

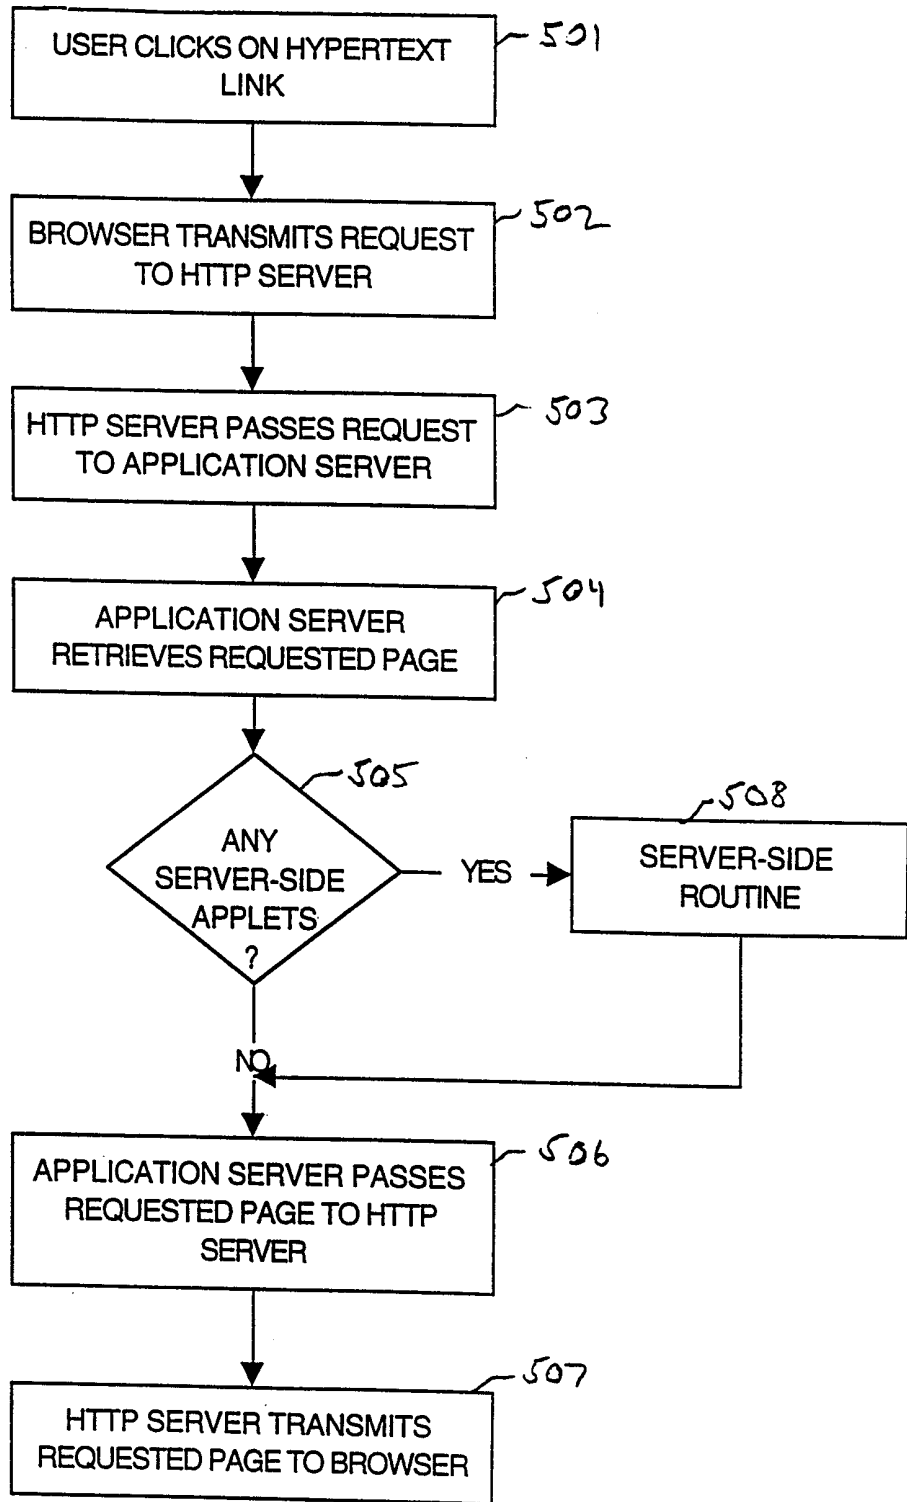


FIG. 5

7/14

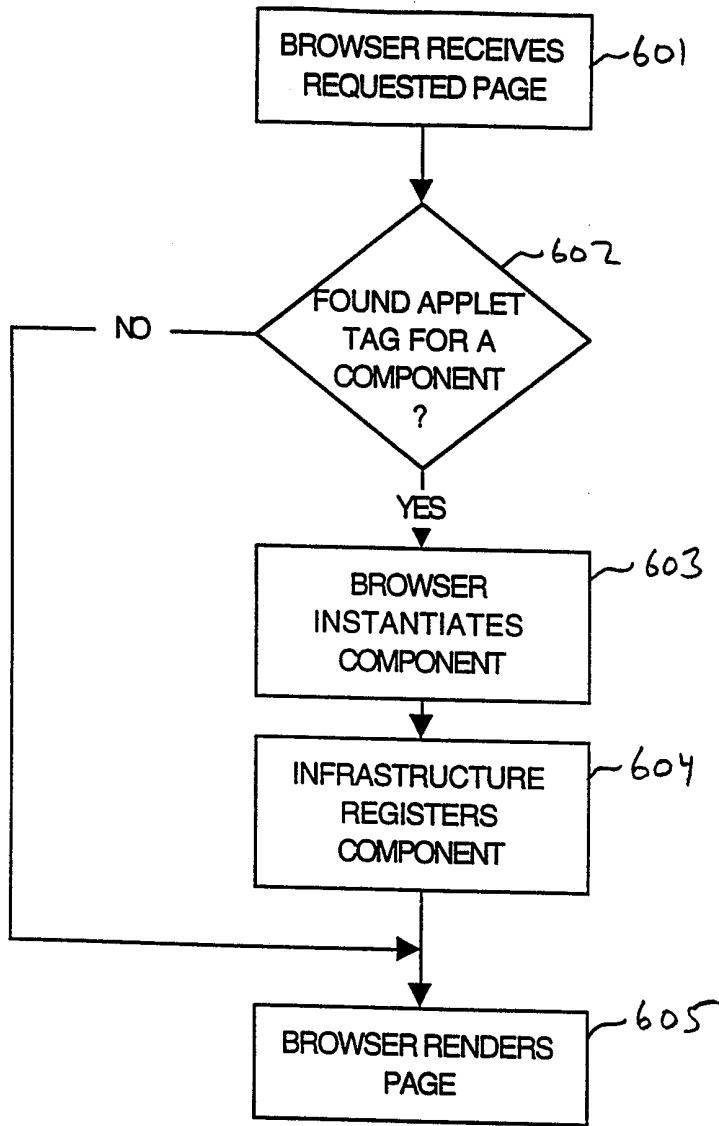


FIG. 6




Fig. 7

9/14

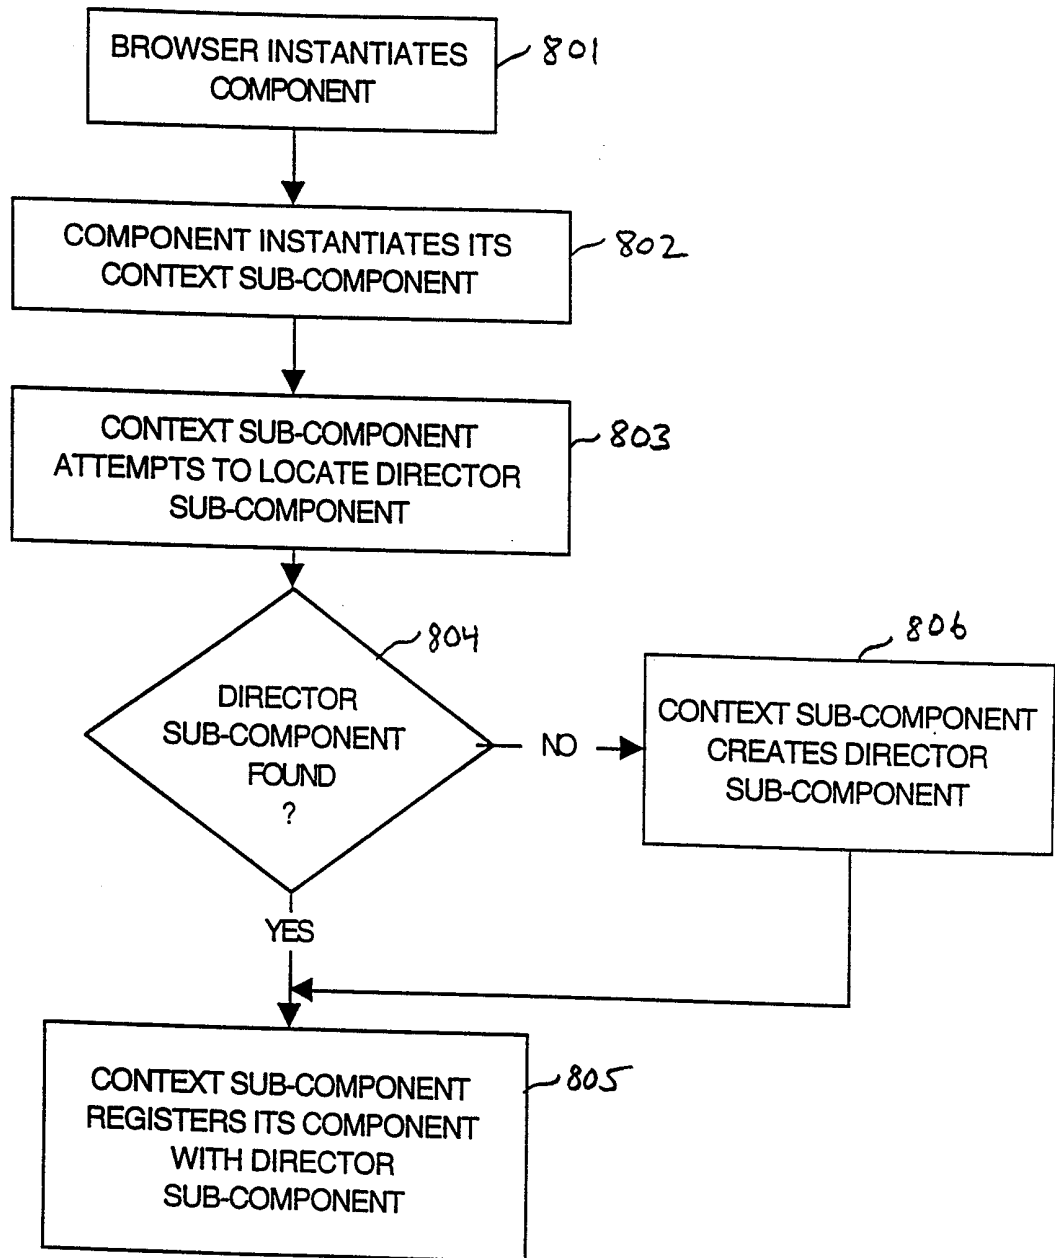


FIG. 8A

10/14




FIG. 8B

11/14




FIG. 9




FIG. 10

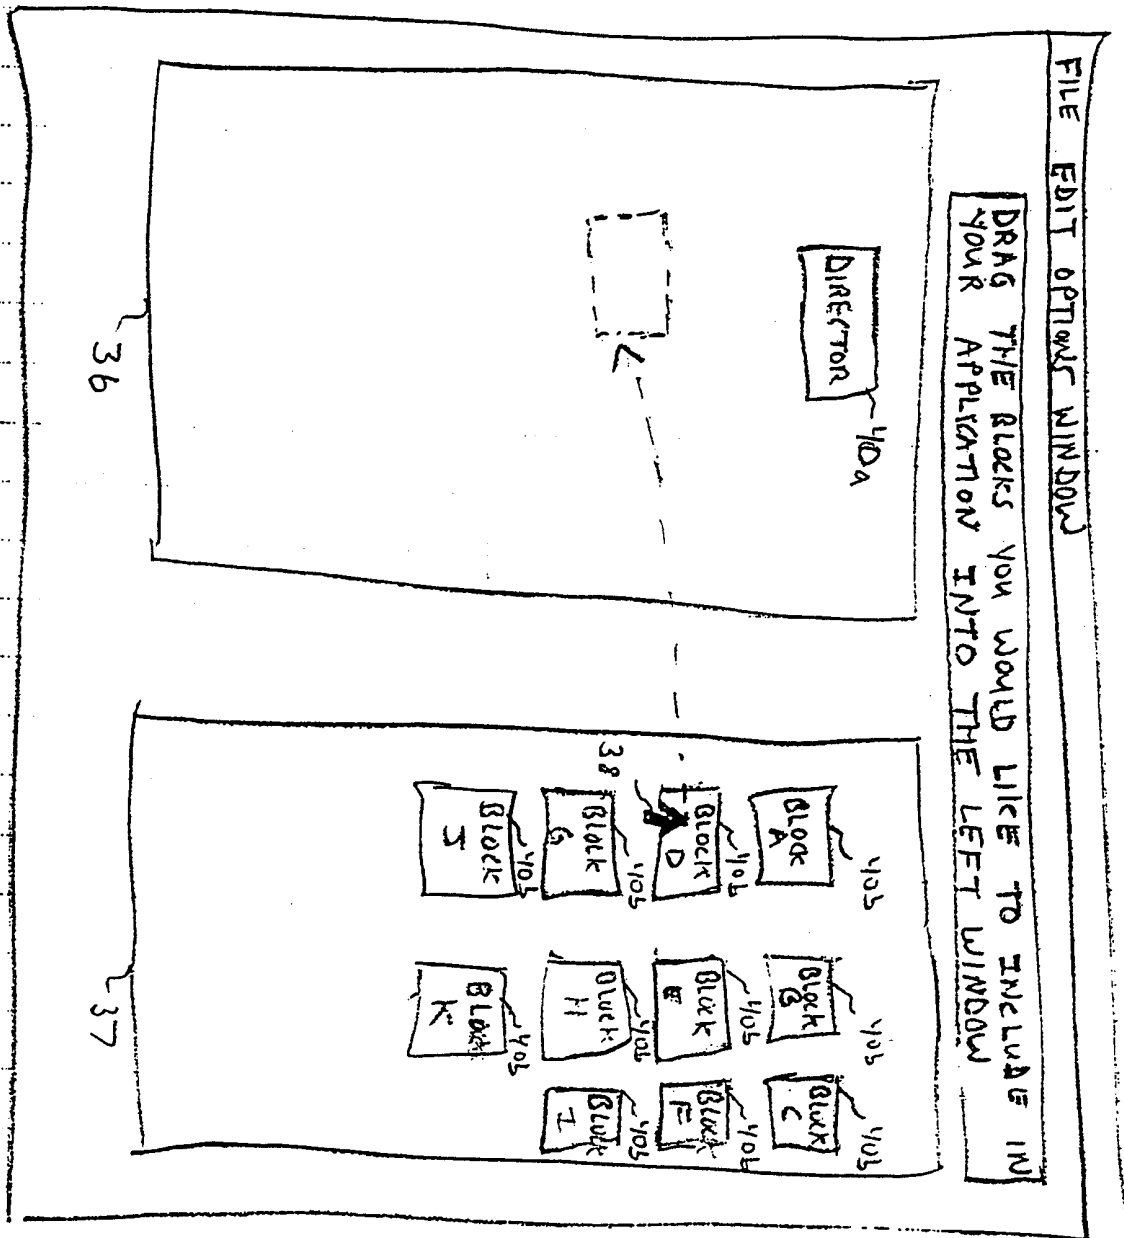


FIG. 11

← 35

50




FIG. 12

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US98/25275

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) :G06F 09/42, 15/16, 17/30
US CL :395/701, 395/685, 395/682

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)

U.S. : 395/701, 395/685, 395/682

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
SOFTWARE DEVELOPMENT 1999 RESOURCE GUIDE

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
APS, STN, ACM

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X,P	USA 5,838,906 A (DOYLE ET AL.) 17 NOVEMBER 1998, ABSTRACT, BACKGROUND, SUMMARY, COLS. 9-11 AND 14-15.	1-59
Y	USA 5,457,797 A (BUTTERWORTH ET AL.) 10 OCTOBER 1995, ABSTRACT, BACKGROUND, SUMMARY, AND COLS. 7-11.	1-59
A	NEIL JENKINS ET AL., CLIENT/ SERVER UNLEASHED, SEPTEMBER 1996, 395-405.	1,7,15,20, 24,28,36, 45,53

Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:	*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
A document defining the general state of the art which is not considered to be of particular relevance	*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
E earlier document published on or after the international filing date	*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
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)	*&* document member of the same patent family
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	

Date of the actual completion of the international search 27 JANUARY 1999	Date of mailing of the international search report 03 MAR 1999
Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703) 305-3230	Authorized officer JOHN CHAVIS <i>Joni Hill</i> Telephone No. (703) 305-9600