

(19) World Intellectual Property Organization
International Bureau



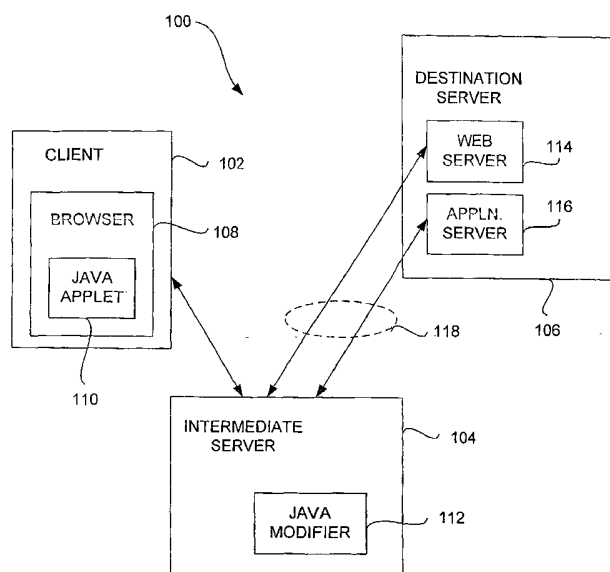
(43) International Publication Date
18 December 2003 (18.12.2003)

PCT

(10) International Publication Number
WO 03/105010 A1

- (51) International Patent Classification⁷: **G06F 15/16**
- (21) International Application Number: PCT/US03/17934
- (22) International Filing Date: 5 June 2003 (05.06.2003)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/387,146 6 June 2002 (06.06.2002) US
- (71) Applicant (*for all designated States except US*): **NEO-TERIS, INC.** [US/US]; 161 East Evelyn Avenue, Mountain View, CA 94041 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (*for US only*): **TOCK, Theron** [US/US]; 620 Palo Alto Avenue, Mountain View, CA 94041 (US). **XIA, Zeqing** [CN/US]; 1399 Chelsea Drive, Los Altos, CA 94024 (US).
- (74) Agent: **THOMAS, C., Douglass**; BEYER WEAVER & THOMAS, LLP, 2030 Addison Street, 7th Floor, P.O. Box 778, Berkeley, CA 94704 (US).
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NI, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**
- with international search report
 - before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments
- For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

(54) Title: METHOD AND SYSTEM FOR PROVIDING SECURE ACCESS TO PRIVATE NETWORKS



(57) Abstract: Improved approaches for providing secure remote access to resources maintained on private networks are disclosed. According to one aspect, predetermined elements, such as applets, can be modified to redirect all communications to and from an application server through an intermediate server. The intermediate server in turn communicates with the application servers. According to another aspect, a communication framework can be provided to funnel communication between an applet and a server through a communication layer so as to provide managed and/or secured communications there between.

WO 03/105010 A1

METHOD AND SYSTEM FOR PROVIDING SECURE ACCESS TO PRIVATE NETWORKS

BACKGROUND OF THE INVENTION

Field of the Invention

[0001] The present invention relates to client-server computing and, more particularly, to client-server computing for securely accessing resources over a network.

Description of the Related Art

[0002] Network browsers (browser applications), such as Netscape Navigator or Microsoft Explorer, allow users of client machines to request and retrieve resources from remotely located server machines via the Internet. These network browsers can display or render HyperText Markup Language (HTML) documents provided by the remotely located server machines. Additionally, browsers are able to execute script programs embedded in the HTML documents to provide some local functionality.

[0003] Further, applets (e.g., Java™ applets) can also be embedded in the HTML documents. In such case, the browser will fetch the bytecode for the applet from a web server by issuing HTTPS requests to get the appropriate class and/or archive files for the applet. The received bytecode is then loaded into a virtual machine (e.g., Java Virtual Machine). During runtime, the applet typically communicates with an application server over a secure connection, such as HTTPS or socket connections. Further, in the case of Java, the Java Sandbox operates to restrict the applet from communicating with a network domain (host) other than the network domain from which the applet was obtained.

[0004] Conventionally, network browsers are used to access public networks, such as the Internet. Private networks are normally protected by firewalls so that network browsers residing on computing machines outside the private network are not able to gain access to any resources on the private network.

[0005] While firewalls are effective at protecting against external access to private networks, there is often the need for external persons or businesses to gain at least limited access to the private networks of other persons or businesses. For example, a supplier of parts to a business customer may be able to better serve their

business customer by having access to information (e.g., inventory levels or orders) maintained on the private network of the business customer. One conventional approach is to allow the supplier's machine to access the private network through the firewall via a public network. This provides a "hole" in the firewall that seriously compromises the security of the private network. Hence, this conventional approach is normally not permitted if security is an important concern. Another conventional approach is to establish a Virtual Private Network (VPN) with the supplier's machine. Here, the supplier's machine is also able to access the private network through the public network and the firewall, but all data transmissions are encrypted. Some firewalls support VPNs and protocols providing encrypted communications, such as Point-to-Point Tunneling Protocol (PPTP). While VPNs offer remote secure access, they are difficult to arrange, configure and manage. Each VPN must also be provided for each external person or business given access to the private network. Still further, VPNs are costly and each VPN provides some security exposure to the entire private network.

[0006] Thus, there is a need for improved approaches to providing secure remote access to resources maintained on private networks.

SUMMARY OF THE INVENTION

[0007] Broadly speaking, the invention pertains to improved approaches for providing secure remote access to resources maintained on private networks.

[0008] According to one aspect of the invention, predetermined elements, such as applets, can be modified to redirect all communications to and from an application server through an intermediate server. The intermediate server in turn communicates with the application servers. Often the applications are secured by a firewall, but the intermediate server is trusted enough to gain access to the application servers, thereby indirectly allowing the applets to communicate with the application servers via the intermediate server.

[0009] According to another aspect of the invention, a communication framework can be provided to funnel communication between an applet and a server through a communication layer so as to provide managed and/or secured communications there between.

[0010] The invention can be implemented in numerous ways, including as a system, method, device, and a computer readable medium. Several embodiments of the invention are discussed below.

[0011] As a method for modifying a markup language page to redirect resource requests to an intermediate server, one embodiment of the invention includes at least the acts of: identifying, within the markup language page, a predetermined element that includes at least a first network address; and modifying the first network address within the predetermined element of the markup language page to a second network address that pertains to the intermediate server.

[0012] As a method for processing resource requests provided to an intermediate server from a client via a computer network, one embodiment of the invention includes at least the acts of: receiving a resource request for a particular resource, the resource request being provided to the intermediate server from the client via the computer network; extracting a destination server from the resource request; requesting the particular resource from the destination server; receiving the particular resource from the destination server; modifying the particular resource to redirect internal resource requests to the intermediate server; sending the modified particular resource to the client; receiving an applet code request for an applet identified within the modified particular resource; requesting applet code for the applet from a remote server via the computer network; receiving the applet code from the remote server in response to the requesting of the applet code; modifying the applet code to redirect its external communications through the intermediate server; and sending the modified applet code to the client.

[0013] As a system for communicating between a client and a server, one embodiment of the invention includes at least: a communication layer at a client, the communication layer transforming one or more socket connections into a pair of unidirectional secure URL connections; an applet operating at the client to perform operations and to create at least one socket connection with the communication layer; and a server operatively connected with the pair of unidirectional secure URL connections, the server communicating with the applet via the pair of unidirectional secure URL connections provided by the communication layer.

[0014] As a system for communicating between a client and a server, one embodiment of the invention includes at least: a plurality of browser applications, at least a plurality of the browser applications utilizing at least one operating applet and a

communication layer, the applet operating to perform operations and to create at least one socket connection with the communication layer, and the communication layer for each of the browsers operates to transform the socket connections into a pair of unidirectional URL connections; and a server operatively connected with the pair of unidirectional URL connections associated with the communication layer associated with each of the plurality of browser applications, the server communicating with the at least one operating applet of the plurality of browser applications via the pair of unidirectional URL connections provided by the communication layer corresponding thereto.

[0015] Other aspects and advantages of the invention will become apparent from the following detailed description taken in conjunction with the accompanying drawings which illustrate, by way of example, the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016] The invention will be readily understood by the following detailed description in conjunction with the accompanying drawings, wherein like reference numerals designate like structural elements, and in which:

FIG. 1 is a block diagram of a computing environment according to one embodiment of the invention.

FIG. 2 is a flow diagram of web resource request processing according to one embodiment of the invention.

FIG. 3 is a flow diagram of applet request processing according to one embodiment of the invention.

FIG. 4 is a flow diagram of response modification processing according to one embodiment of the invention.

FIG. 5 is a flow diagram of bytecode modification processing according to one embodiment of the invention.

FIG. 6 is a block diagram of a communication framework according to one embodiment of the invention.

FIG. 7 is a block diagram of a communication framework according to another embodiment of the invention.

DETAILED DESCRIPTION OF THE INVENTION

[0017] According to one aspect of the invention, applets are modified to redirect all communications to and from an application server through an intermediate server. The intermediate server in turn communicates with the application servers. Often the applications are secured by a firewall, but the intermediate server is trusted enough to gain access to the application servers, thereby indirectly allowing the applets to communicate with the application servers via the intermediate server.

[0018] FIG. 1 is a block diagram of a computing environment 100 according to one embodiment of the invention. The computing environment 100 includes a client 102, an intermediate server 104, and a destination server 106. The client 102 pertains to a computing device, such as a personal computer, desktop computer or Personal Digital Assistant (PDA) for example. A browser 108 is a program that operates on the client 102 to send, receive and display resources residing on remote servers. A Java applet 110 is another application that is operated within the browser 108.

[0019] The intermediate server 104 includes a Java modifier 112 that operates to modify a Java applet before supplying the Java applet to the client 102. The modified Java applet, once provided to the client 102, becomes the Java applet 110. The destination server 106 includes a web server 114 and an application server 116. The intermediate server 104 couples to the client 102 through a first network (not shown), and the intermediate server 104 couples to the destination server 106 through the first network or a second network (not shown). The first and second networks can include wired and wireless components, and can comprise the Internet, Local Area Networks, Wide Area Networks, etc. Typically, although not necessarily, the coupling of the intermediate server 104 to the destination server 106 is achieved through a firewall 118. The firewall 118 serves to restrict the ability for external computing devices to gain access to the destination server 106.

[0020] The general operation of the computing environment 100 can be as follows. Typically, a user would interact with the browser 108 operating at the client 102. Such interaction would cause the browser 108 to request a web resource (e.g., HTML page) from the destination server 106. However, the web resource request would be intercepted by the intermediate server 104. The intermediate server 104 would then produce a reformatted web resource request from the web resource request received from the browser 108. Thereafter, the reformatted web resource

request would be sent (e.g., over the second network) to the destination server 106 (e.g., the web server 114). The destination server 106 would then retrieve the requested web resource and return it to the intermediate server 104. At this point, the intermediate server 104 can operate to modify the web resource to facilitate redirection of subsequent resource requests internal to the web resource to the intermediate server 104. In other words, the web resource is modified such that any subsequent requests from the requested web resource are directed initially to the intermediate server 104 (as opposed to an appropriate destination server). In particular, a Java modifier 112 at the intermediate server 104 can modify links or attribute values within the web resource. Among other things, the attributes or values being modified within the web resource can include those attributes or values found within Java elements contained within the web resource. The modified web resource is then sent from the intermediate server 104 to the browser 108 via the client 102. Then, when the modified web resource is evaluated (e.g., displayed) by the browser 108, the modified web resource includes the Java applet 110. At this point, the browser 108 will send to the intermediate server 104 a request for the code (i.e., bytecode) for the Java applet 110 to the intermediate server 104. Upon receiving the request for the code for the Java applet, the intermediate server 104 reformats the request. The reformatted request is then sent to the appropriate destination server 106. At the destination server 106, the request code is retrieved and sent to the intermediate server 104. The retrieved code is then modified by the Java modifier 112. Here, the retrieved code (for the Java applet 110) is modified such that communications between the client 102 and the destination server 106 are redirected to the intermediate server 104 in a secure fashion. Despite the redirection, the code for the Java applet 110 is able to communicate with the destination server 106 through the intermediate server 104, even though the code for the Java applet 110 was originally designed to directly communicate with the destination server 106.

[0021] Although FIG. 1 illustrates only a single destination server and a single client, it should be understood that an intermediate server can support many clients and can facilitate access to many destination servers.

[0022] FIG. 2 is a flow diagram of web resource request processing 200 according to one embodiment of the invention. The web resource request

processing 200 is, for example, performed at an intermediate server, such as the intermediate server 104 illustrated in FIG. 1, for example.

[0023] The web resource request processing 200 begins with a decision 202 that determines whether a web resource request has been received. Here, the web resource request processing 200 awaits the receipt of a web resource request from the client. Hence, when the decision 202 determines that a web resource request has not been received, then the web resource request processing 200 awaits such a request. Once the decision 202 determines that a web resource request has been received, a destination server is extracted 204 from the web resource request. Here, the web resource request is originally destined for the intermediate server because such a redirection has been imposed. However, the actual destination server for the web resource request is contained within the web resource request. Hence, the actual destination server can be obtained (extracted 204) from the web resource request being received at the intermediate server from the client.

[0024] After the destination server has been extracted 204 from the web resource request, a replacement request for the web resource can be sent 206 to the destination server. Here, the intermediate server can reformat the original web resource request or otherwise generate the replacement request.

[0025] Next, a decision 208 determines whether a response to the replacement request has been received from the destination server. When the decision 208 determines that a response has not yet been received, then the web resource request processing 200 awaits such a response. Once the decision 208 determines that a response has been received from the destination server, the response is modified 210 to facilitate redirection to the intermediate server. According to one embodiment, the modification to the response involves altering complete or partial Universal Resource Locators (URLs) within the response, altering attribute values, or inserting additional elements. After the response has been modified 210, the modified response is returned 212 to the requestor. In one embodiment, the requestor is a browser operating on the client. In another embodiment, the requestor is the client or the user of the browser or client. Following the operation 212, the web resource request processing 200 is complete and ends.

[0026] Once the modified response is received at the client, the browser at the client typically displays the modified response. In addition, when the modified response includes internal resource requests, such as for images and applets, the

browser makes subsequent requests for such additional resources. The processing of a subsequent request for an applet is described below with respect to FIG. 3.

[0027] FIG. 3 is a flow diagram of applet request processing 300 according to one embodiment of the invention. The applet request processing 300 is, for example, performed at an intermediate server, such as the intermediate server 104 illustrated in FIG. 1, for example.

[0028] The applet request processing 300 begins with a decision 302 that determines whether an applet request has been received. When the decision 302 determines that an applet request has not yet been received, the applet request processing 300 awaits such a request. Once the decision 302 determines that an applet request has been received, then bytecode for the applet is requested 304. Typically, the bytecode for the applet will reside on the same destination server from which the web resource (including the java applet) was obtained. Hence, the bytecode for the applet will be requested 304 from the destination server.

[0029] Next, a decision 306 determines whether the bytecode has been received from the destination server. When the decision 306 determines that the bytecode has not yet been received at the intermediate server, then the applet request processing 300 awaits its arrival. Once the decision 306 determines that the bytecode has been received, then the bytecode is modified 308. The modifications made to the bytecode facilitate redirection of communications between the client and the destination server through the intermediate server. After the bytecode has been modified 308, the modified bytecode is sent to the requestor. In one embodiment, the requestor is a browser operating on the client. In another embodiment, the requestor is the client or the user of the browser or client. Following the operation 310, the applet request processing 300 is complete and ends.

[0030] FIG. 4 is a flow diagram of response modification processing 400 according to one embodiment of the invention. The response modification processing 400 is, for example, processing performed by the operation 210 illustrated in FIG. 2. In FIG. 4, the response being modified is assumed to be an HTML page. However, it should be understood that the response being modified could be some other sort of web resource, including any markup language document.

[0031] The response modification processing 400 initially scans 402 an HTML page for <applet> or <object> elements. Here, the <object> being scanned for is

typically a Java object. The <applet> or <object> elements specify embedded Java applets within the HTML page. A decision 404 then determines whether an <applet> or <object> element has been found. When the decision 404 determines that such objects have not been found during the scanning 402 of the HTML page, then a decision 406 determines whether scanning has been completed. When the decision 406 determines that scanning has not been completed, then the modify response processing 400 returns to repeat the operation 402 and subsequent operations. Alternatively, when the decision 406 determines that scanning has been completed, the response modification processing 400 is complete and ends.

[0032] On the other hand, when the decision 404 determines that either <applet> or <object> elements have been found, then a decision 408 determines whether the element includes a "codebase" attribute. When the decision 408 determines that the element does not include a "codebase" attribute, then a default "codebase" attribute is inserted 410 into the element. Alternatively, when the element does include a "codebase" attribute, the operation 410 is bypassed.

[0033] Subsequently, URLs associated with certain attributes of the element are modified 412. Here, the URLs being modified 412 can include complete or partial URLs. For example, the attributes of the element being modified can include a "codebase" attribute as well as various parameters. Further, additional parameters can also be inserted 414 into the element. In particular, additional parameters can be inserted into the element to store original URLs of the document and the codebase. Following the operation 414, the response modification processing 400 is complete and ends with the response having been modified.

[0034] According to one implementation, the operations 412 and 414 of the applet request processing 400 are as follows. Assume that the network address (e.g., URL) for the intermediate server is "http://www.danastreet.com." One of the URLs being modified is a codebase attribute that identified the location of the bytecode for the applet. The codebase attribute is either an absolute (complete) URL or a relative (partial) URL. In one implementation, if the codebase specifies an absolute URL, it is translated to a URL pointing to the intermediate server, with the original host and path information embedded.

[0035] For example, an original URL of "http://xyz.com:123/abc" could be translated to https://www.danastreet.com/abc/DanaInfo=xyz.com,Port=123,CT=java" where "DanaInfo" identifies the original host, "Port" identifies the original port, and

"CT" indicates the content type (e.g., Java applet). Alternatively, if the codebase specifies only a relative URL (which uses the base of the HTML page), e.g., "abc/xyz", the intermediate server will first construct the absolute URL by using the URL of the HTML page (or specified in a <base> tag) and then will transform the resulting URL using the above approach. Note that when the default "codebase" attribute is to be inserted 410, the "DanaInfo" is able to be preserved, and thus not lost when files from the same directory are fetched. Another of the URLs within <applet>...</applet> tags being modified are the parameters for "codebase value" and "cabbase value." More particularly, the value specified for <param name=codebase value=...> will be transformed using the above approach. The value specified for <param name=cabbase value=...> will be transformed using the above approach if it is an absolute URL. In addition, two new applet parameters "neoteris-doc-base" and "neoteris-code-base" are injected before </applet> to store the values of the original document URL and the original codebase URL as they were prior to the above transformations.

[0036] A representative example of this implementation of the operations 412 and 414 of the applet request processing 400 are as follows. Assume that an HTML page at URL "http://www.acmegizmo.com/myapplet.html" contains an applet:

```
<applet codebase="myapplet/classes" archive="applet.jar">
<param name=cabbase value="applet.cab">
</applet>
```

According to the representative example, the applet will be rewritten (transformed) to:

```
<applet
codebase="myapplet/classes,DanaInfo=www.acmegizmo.com,CT=java+" archive="applet.jar">
<param name=cabbase value="applet.cab">
<param name=neoteris-code-base
value=http://www.acmegizmo.com/myapplet/classes>
<param name=neoteris-doc-base
value=http://www.acmegizmo.com/myapplet.html>
</applet>
```

[0037] FIG. 5 is a flow diagram of bytecode modification processing 500 according to one embodiment of the invention. The bytecode modification

processing 500 is, for example, performed by a java modifier, such as the java modifier 112 illustrated in FIG. 1. In addition, the bytecode modification processing 500 is, for example, suitable for use as the processing performed by the operation 308 illustrated in FIG. 3.

[0038] By modification of the bytecode, the behavior of an applet is modified such that all HTTP and socket based network communications are redirected to an intermediate server over HTTPS connections. The intermediate server manages connection end points with actual application servers and forwards network traffic between Java applets and application servers.

[0039] The bytecode modification processing 500 initially scans 502 bytecode of the Java applet for a class descriptor. Once a class descriptor has been identified in the bytecode, a decision 504 determines whether the class is a final class. When the decision 504 determines that the class is not a final class, then each identified class descriptor is replaced 506 with a corresponding modified class that is a subclass of the identified class.

[0040] Alternatively, when the decision 504 determines that the class is a final class, then a decision 508 determines whether the class has public constructors. When the decision 508 determines that the class does have public constructors, then a decision 510 determines whether object method invocations are present. When the decision 510 determines that object method invocations are present, then each object method invocation is replaced 512 with a corresponding static method call. Following the operation 512, as well as directly following the decision 510 when no object method invocations are present, a decision 514 determines whether object creations are present. When the decision 514 determines that object creations are present, then each object creation is replaced 516 with a corresponding static method. On the other hand, when the decision 508 determines that the class does not have public constructors, then predetermined class strings in the class are replaced 518 with corresponding wrapper classes.

[0041] Following the operations 506, 516 and 518, as well as directly following the decision 514 when object creations are not present, a decision 520 determines whether scanning of the HTML page has been completed. When the decision 520 determines that scanning has not been completed, then the bytecode modification processing 500 returns to repeat the operation 502 and subsequent operations.

Alternatively, when the decision 520 determines that scanning has been completed, the bytecode modification processing 500 is complete and ends.

[0042] Bytecode files for a Java applet are normally stored at an application server within a destination server. The bytecode files are either stored individually (e.g., .class files), or bundled into archive files. The archive files can be in the format of .zip files, .jar files, or .cab files, each using a different compression method and specification. Upon fetching the bytecode files, and uncompressing the archive files if necessary, an intermediate server applies one or more of the following four techniques in modifying each individual bytecode class file. These techniques and examples provided below are described at the Java language level; however, these modification are preferably performed at the bytecode level.

[0043] According to one implementation of a first technique, the operation 506 of the bytecode modification processing 500 is as follows. According to the first technique, when an original class is not final, a new subclass is created and then all objects of the original class are replaced with the new class. For example, whenever a "java.applet.Applet" class is encountered, it is replaced with a "com.neoteris.NeoterisApplet" class, which is a subclass of "java.applet.Applet". Since all applets contain subclasses of "java.applet.Applet", this essentially makes them all subclasses of "com.neoteris.NeoterisApplet".

[0044] A representative Java language level example of this implementation of the first technique is as follows. An initial applet, MyApplet, initially includes

```
import java.applet.*;

public MyApplet extends Applet {
...
}
```

and is effectively modified as shown below to become a subclass of NeoterisApplet.

```
import com.neoteris.NeoterisApplet;

public MyApplet extends NeoterisApplet {
...
}
```

NeoterisApplet modifies the methods of the applet to reverse the effect of rewriting of codebase so that the inquiring of applet information (e.g., codebase) from within the applet's application logic will yield the same result as if there had been no modification to the bytecode.

[0045] The implementation of the first technique further modifies socket interfaces in the bytecode so that socket connections are transformed into HTTPS connections. For example, whenever a "java.net.Socket" class is encountered, it is replaced with "com.neoteris.NeoterisSocket" or subclasses of it. A representative Java language level example of this implementation of the first technique for socket connections is as follows. Assume that the applet originally contained,

```
import java.net.Socket;
import java.io.*;

Socket sock = new Socket("myserver.com", 1234);
InputStream is = sock.getInputStream();
...
```

and then is effectively modified as shown below to alter the nature of the socket connections.

```
import com.neoteris.NeoterisSocket;
import java.io.*;

Socket sock = new NeoterisSocket("myserver.com", 1234);

// will return a NeoterisInputStream object
InputStream is = sock.getInputStream();
...
```

[0046] The implementation of the first technique still further modifies socket interfaces in the bytecode so that datagram sockets are similarly modified. For

example, whenever a "java.net.DatagramSocket" class is encountered, it is replaced with "com.neoteris.NeoterisDatagramSocket" or subclasses of it.

[0047] According to one implementation of a second technique, the operation 512 of the bytecode modification processing 500 is as follows. According to the second technique, object method invocations with the original class are changed to static methods (calls to static methods) so that communications with the intermediate server operate properly. For example, whenever the following method invocations on "java.net.URL" are encountered, they are replaced with the corresponding static methods of class "com.neoteris.NeoterisStatic" (with the same name). These methods are, for example: getHost, getPort, getFile, getProtocol, openStream and openConnection. Further, other methods, namely, java.net.Socket: getInetAddress and java.net.DatagramSocket: getAddress and setAddress, are also transformed to static methods in NeoterisStatic (with the same name), the return values being NeoterisInetAddress.

[0048] According to one implementation of a third technique, the operation 516 of the bytecode modification processing 500 is as follows. According to the third technique, creation of an object class is replaced with a static method. Because of the inability to subclass a Java class that is declared final, it is necessary to replace the bytecode instruction that creates an object with static methods in order to inject additional code to process parameters. For example, the creation of an object of class java.net.URL is replaced with the static method translateURL of class com.neoteris.NeoterisStatic. As another example, the creation of an object of class java.net.DatagramPacket is replaced with the static method translateDatagramPacket of class com.neoteris.NeoterisStatic.

[0049] A representative Java language level example of this implementation of the second and third techniques for inserting static methods is as follows. Assume that the applet originally contained,

```
URL url = new URL("http", "myserver.com", 80,  
                  "cgi/dothis?task=1");  
String host = url.getHost();
```

and then is effectively modified as shown below to inject the static methods.

```
// Will become
// "http://www.danastreet.com/cgi/dothis?task=1,DanaInfo=
// myserver.com,Port=80,CT=txt"
URL url = NeoterisStatic.translateURL("http",
    "myserver.com", 80, "myapplet/dothis?task=1");

// Will reverse the effect of URL translation and return
// host string "myserver.com"
String host = NeoterisStatic.getHost(url);
```

[0050] According to one implementation of a fourth technique, the operation 518 of the bytecode modification processing 500 is as follows. According to the fourth technique, class strings are replaced with a specialized string. In cases where a standard Java class is a final class and has no public constructors, such as `java.net.InetAddress`, all occurrences of the class string, (in both class name and in method descriptors) are replaced with a specialized class that implements communications through an intermediate server. For example, for all occurrences of `java.net.InetAddress`, they are replaced with `com.neoteris.NeoterisInetAddress`. More particularly, this would include strings in the constant (const) pool of the bytecode that represent the class, i.e., "`java/net/InetAddress`", and any appearance of "`java/net/InetAddress`" in a descriptor, such as "`(Ljava/net/InetAddress;String)V`".

[0051] A representative Java language level example of this implementation of the second and fourth techniques for performing string replacements is as follows. Assume that the applet originally contained,

```
public class MyClass {
    public String getInfo(InetAddress addr, int type);
    ...
}

MyClass cl = new MyClass();
Socket sock = new Socket(...);
cl.getInfo(sock.getInetAddress, 2);
```

and then is effectively modified to the following.


```

public class MyClass {
    public String getInfo(NeoterisInetAddress addr, int
type);
    ...
}
MyClass cl = new MyClass();
Socket sock = new NeoterisSocket(...);
cl.getInfo(NeoterisStatic.getInetAddress(sock), 2);

```

[0052] According to one implementation of a fifth technique, although not shown in FIG. 5, somewhere in the scanning operation of the bytecode modification processing 500, static class methods of a Java class are replaced with static class methods of a specialized class (e.g., NeoterisStatic). A representative Java language level example of this implementation of the fifth technique for static class methods is as follows. Assume that the applet originally contained,

```

JSObject object = JSObject.getWindow(this_applet);

```

and then is effectively modified as shown below

```

JSObject object = NeoterisStatic.getJSWindow(this_applet);

```

[0053] According to another aspect of the invention, a communication framework is provided to funnel communication between an applet and a server through a communication layer so as to provide managed and/or secured communications there between. In one embodiment, socket connections from one or more applets are funneled into HTTPS connections, namely, a pair of unidirectional HTTPS connections.

[0054] FIG. 6 is a block diagram of a communication framework 600 according to one embodiment of the invention. The communication framework 600 provides for communications between applets operating at a client with destination servers (or remote content servers). The communication framework 600 shown in FIG. 6 includes an applet 602, an applet communication layer 604 and an intermediate

server 606. The general communications scheme is that the applet 602 sends outgoing communications through the applet communication layer 604 and then onto the intermediate server 606 before being delivered to an appropriate destination server. Typically, the communications would be delivered to an application server within the appropriate destination server. Also, communications incoming to the applet 602 are initially received at the intermediary server 606 (from the destination server or an application server thereof). The incoming communications are then directed by the intermediate server to the applet communication layer 604. The applet communication layer 604 then directs the incoming communications to the applet 602.

[0055] The applet 602 includes a plurality of communication sockets (socket connections), namely, socket A 608, socket B 610 and socket C 612. Each socket effectively provides a communication link, channel or connection to an eventual destination server. The sockets 608-612 provide communications in both incoming and outgoing directions. Normally, the communication is achieved through transmission of packets of data. The outgoing packets from the sockets 608-612 are supplied to an outgoing queue 614 of the applet communication layer 604. In other words, the outgoing queue 614 queues the packets being sent from the applet 602 to the applet communication layer 604. A sender thread 616 interacts with the outgoing queue 614 to send those packets residing in the outgoing queue 614 to the intermediate server 606. The communications between the sender thread 616 of the applet communication layer 604 and the intermediate server 606 are preferably achieved over a secure connection. One such secure connection is a HTTPS connection. Once the packets arrive at the intermediate server 606, dispatching logic 618 determines which of a plurality of intermediate sockets the packets should be directed to. As depicted in FIG. 6, the intermediate server 606 is shown including intermediate socket A 620, intermediate socket B 622 and intermediate socket C 624. Each of the intermediate sockets 620-624 couples to and provides a socket connection with a destination server (or application server therein), normally a different destination server for each of the intermediate sockets 620-624.

[0056] Incoming communications are directed from one of the destination servers to one of the intermediate sockets 620-624 of the intermediate server 606. Normally, the communication is achieved through transmission of packets of data. The communication of these packets from the intermediate server 606 to a reader thread

626 of the applet communication layer 604 is achieved over a secure connection, namely, a HTTPS connection. The reader thread 626 then determines which of a plurality of buffers the incoming packets should be directed to. Namely, the applet communication layer 604 includes buffers 628-632. Each of the buffers 628-632 corresponds to one of the sockets 608-612 within the applet 602. Hence, packets destined for the socket A 608 of the applet 602 are directed to the buffer 628 by the reader thread 626. Similarly, packets directed to the socket B 610 are directed to the buffer 630 by the reader thread 626, and packets destined for the socket C 612 are directed to the buffer 632 by the reader thread 626. When the sockets of the applet 602 read from the socket's input stream, it retrieves any packets for the socket that are in the corresponding buffer. Once a socket of the applet 602 is destroyed, the corresponding input and output stream objects, along with underlying data structures (e.g., buffers), are destroyed.

[0057] The connections between the applet communication layer 604 and the intermediary server 606 are provided by HTTPS. Namely, one HTTPS connection for outgoing packets, and another HTTPS connection for incoming packets. Each of these HTTPS connections is half-duplexed, meaning that it supports communications in one direction. Additionally, the outgoing HTTPS connection from the applet communication layer 604 is a transient connection. The transient connection is such that the sender thread 616 establishes the connection while packets are stored in the outgoing queue 614 and ready to be delivered. Once the sender thread 616 has sent all the available packets from the outgoing queue 614, the outgoing HTTPS connection is closed. On the other hand, the incoming HTTPS connection to the applet communication layer 604 (namely, the reader thread 626) is maintained as a persistent connection (at least until a time-out occurs). In other words, to timely deliver incoming packets to the appropriate sockets within the applet 602, the reader thread 626 constantly monitors the incoming HTTPS connection to discover incoming data packets.

[0058] By using the HTTPS connections between the applet communication layer 604 and the intermediate server 606, the communications there between are secured. Also, the applet communication layer 604 allows the pair of HTTPS connections to support numerous sockets (socket connections) within the applet 602. Therefore, the limitation of Java applets that only a maximum of two (2) URL

connections be permitted is able to be satisfied, yet the applet or applets are able to have many more active sockets than two.

[0059] FIG. 7 is a block diagram of a communication framework 700 according to another embodiment of the invention. The communication framework 700 illustrates use of a plurality of browsers sharing an intermediate server and then coupling to various different destination servers through the shared intermediate server. Applets operating on these browsers are able to communicate over secure connections to the intermediate server which directs communications to and from the appropriate destination server (or application server therein).

[0060] The communication framework 700 illustrated in FIG. 7 includes a browser A 702 and a browser B 704. The browser A 702 includes an applet A 706. The applet A 706 has socket connections 708 and 710. These socket connections 708 and 710 communicate through an applet communication layer 712 with an intermediate server 714. In one embodiment, the applet communication layer 712 is designed as shown in FIG. 6. Typically, the applet communication layer 712 communicates with the intermediate server 714 over a pair of HTTPS connections. Each of the HTTPS connections provides secure communications in a single direction (i.e., unidirectional).

[0061] The browser B 704 includes an applet B1 716 and an applet B2 718. The applet B1 716 includes socket connections 720 and 722, and the applet B2 718 includes socket connection 724. The socket connections 720-724 communicate with the intermediate server 714 through an applet communication layer 726. Here, the applet communication layer 726 is able to support socket connections from one or more applets operating on the browser B 704. The applet communication layer 726 couples to the intermediate server 714 through a pair of HTTPS connections. Each of the HTTPS connections provide secure communications in a single direction (i.e., unidirectional).

[0062] The intermediate server 714 includes a plurality of intermediate socket connections. Namely, the intermediate server 714 includes intermediate socket connection-1 728, intermediate socket connection-2 730, intermediate socket connection-3 732, intermediate socket connection-4 734, and intermediate socket connection-5 736. The intermediate socket connection-1 728 provides a connection with a destination server-1 738. The intermediate socket connection-2 730 provides a connection with a destination server-2 740. The intermediate socket connection-3

732 provides a connection to a destination server-3 742. The intermediate socket connection-4 734 provides a connection to a destination server-4 744. The intermediate socket connection-5 736 provides a connection to a destination server- 5 746.

[0063] Additionally, the socket connections 708, 710 and 720-724 in the applets 716 and 718 are associated with the intermediate socket connections 728-736 in the intermediate server 714. Namely, the applet A 706 forms the socket connection 708 to communicate with the destination server-1 738. However, the delivery of packets for communications between the socket connection 708 and the destinations server-1 738 is achieved through the applet communication layer 712 and the intermediate server 714. In particular, the socket connection 708 sends packets to the applet communication layer 712 that forwards the packets over the HTTPS connection to the intermediate server 714. The intermediate server 714 recognizes that the incoming packets originated from the socket connection 708 of the applet A 706 and thus directs the packets to the intermediate socket connection-1 728. The intermediate server 714 then delivers the packets from the intermediate socket connection-1 728 to the destination server-1 738. Similarly, packets being sent from the socket connection 710 of the applet A 706 are directed through the applet communication layer 712 to the intermediate socket connection-2 730 of the intermediate server 714 and then onto the destination server-2 740. Likewise, packets being sent by the socket connections 720, 722, and 724 from the applets 716 and 718 are directed by the applet communication layer 726 to the intermediate socket connections 732, 734 and 736, respectively, within the intermediate server 714. The intermediate socket connections 732, 734 and 736 then respectively deliver the packets to the corresponding destination servers 742, 744 and 746.

[0064] In one embodiment, the packets being transmitted include a header that identifies the socket connection and the browser hosting the socket connection. Using such information from the header, the intermediate server can determine the appropriate intermediate socket connection within the intermediate server 714 that is to handle the packets.

[0065] Optionally, the intermediate server can impose security on which destination servers the sockets for the applets are able to communicate with. Hence, the sockets of an applet can also be assigned a security identifier. The security identifier can be utilized to restrict the network locations that the socket is

able to communicate with. For example, if the HTML page including an applet originated from xyz.com, when subsequent packets from the socket connection are sent to the intermediate server, the intermediate server can check the security identifier to determine which destination server the applet is permitted to communicate with. Hence, in one embodiment, the intermediate server can maintain a table of security identifiers and permitted network domains. For example, if the security identifier was associated with xyz.com, then packets to be sent to a domain jkl.com could be refused by the intermediate server. The security identifier can thus be used to enforce the network security described in the Java security specifications. In one implementation, the security can be implemented by bytecode modification discussed above, wherein the security identifier for each applet can be inserted into the bytecode.

[0066] The communication frameworks 600, 700 discussed above in FIGs. 6 and 7 make use of a pair of HTTPS connections between an applet communication layer (client) and an intermediate server (server). As a result, the data packets (or messages) from various browsers can be multiplexed and sent over one of the HTTPS connections, and the data packets (or messages) destined for various browsers can be commonly sent over another of the HTTPS connections and then delivered to the appropriate browser.

[0067] To assist the exchange of data packets (messages) flow control messages can also be sent between the applet communication layer and the intermediate server for flow control. If the intermediate server is sending too much data, the applet communication layer can inform the intermediate server to stop sending the data. This prevents the applet communication layer from being forced to buffer more than a limited amount of data. After at least a significant portion of the data being buffered in the buffers 628-632 of the applet communication layer 604 has been consumed by the corresponding applet(s), the applet communication layer can instruct the intermediate server to resume sending more data. This flow control can be useful because the sockets are being multiplexed across a single HTTPS connection in each direction.

[0068] Additionally, acknowledgement messages can also go back and forth between the applet communication layer and the intermediate server so that the intermediate server knows what data has actually been received by the applet communication layer. The use of acknowledgement messages can be beneficial as

sometimes the long-lived HTTPS connection (the connection that the applet communication layer uses to read data from the intermediate server) can be dropped by an intermediate proxy if the connection is idle or open for an extended duration. To recover from this situation the intermediate server is capable of retransmitting any data that the applet communication layer did not receive because the intermediate server knows from the acknowledgements what data the applet communication layer did or did not receive.

[0069] It should be recognized that the amount of overhead involved in the dual-channel implementation (provided by the pair of HTTPS connections) can be significant. For example, the overhead can include acknowledgements, flow control messages, and the HTTP overhead for each transient HTTP request (sending data from client to server). Note that the browser can also perform HTTP keep-alives, so the underlying TCP connection from browser (namely, the applet communication layer) to intermediate server is kept open. The transient HTTP requests from client to server are thus able to avoid the expensive TCP handshake on every request.

[0070] The dual-channel implementation works in situations where the applet must talk through a proxy to reach the intermediate server. However in situations where the applet can connect directly to the intermediate server, the HTTP and acknowledgement/flow-control overhead can be avoided if the applet communication layer opens an SSL connection directly to the intermediate server. By opening a new TCP connection for each client socket connection (e.g., NeoterisSocket) there is no longer any multiplexing of connections. Thus if the applet stops reading from a particular client socket connection, the data on the corresponding TCP connection will be unread and, due to the flow control built into TCP, the intermediate server will be unable to send any more data. Once the applet reads from data from the client socket connection, a read from the TCP connection can be performed, which permits the intermediate server to send more data. Here, there is no proxy in the middle that could drop connections, so the acknowledgements are not necessary.

[0071] Moreover, because opening a new SSL connection is expensive, it is desirable to keep an SSL connection open for as long as possible. If the applet opens a connection sends a small amount of data and then closes the connection, it is desirable to keep the underlying TCP and SSL connection open. This can be achieved by embedding special control messages in the data stream that indicate when the connection is being closed by the application. The TCP and SSL

connections are kept around for a period of time (e.g., 1-5 minutes), so that if a new connection is needed the previous one can be reused. Here, the logic of the applet then is to first try a direct connection to the intermediate server. If that fails, or times out, the applet falls back to the "proxy friendly" method of communication that uses the pair of HTTPS connections.

[0072] Although the above-described embodiments refer to the use of a single intermediary server within an information retrieval system, it should be recognized that the information retrieval system can also include a plurality of intermediary servers. The various intermediary servers can individually receive requests from client machines and forward them to the appropriate servers and return responses back through the intermediary server to the client machine. By having multiple servers, not only can additional processing power be obtained, but load balancing, fault tolerance and localization issues can also be addressed.

[0073] The various aspects, features, embodiments or implementations of the invention described above can be used alone or in various combinations.

[0074] The invention is preferably implemented in software, but can be implemented in hardware or a combination of hardware and software. The invention can also be embodied as computer readable code on a computer readable medium. The computer readable medium is any data storage device that can store data which can thereafter be read by a computer system. Examples of the computer readable medium include read-only memory, random-access memory, CD-ROMs, DVDs, magnetic tape, optical data storage devices, and carrier waves. The computer readable medium can also be distributed over network-coupled computer systems so that the computer readable code is stored and executed in a distributed fashion.

[0075] The many features and advantages of the present invention are apparent from the written description and, thus, it is intended by the appended claims to cover all such features and advantages of the invention. Further, since numerous modifications and changes will readily occur to those skilled in the art, it is not desired to limit the invention to the exact construction and operation as illustrated and described. Hence, all suitable modifications and equivalents may be resorted to as falling within the scope of the invention.

What is claimed is:

CLAIMS

1. A method for modifying a markup language page to redirect resource requests to an intermediate server, said method comprising the acts of:
 - 5 identifying, within the markup language page, a predetermined element that includes at least a first network address; and
 - modifying the first network address within the predetermined element of the markup language page to a second network address that pertains to the intermediate server.
- 10 2. A method as recited in claim 1, wherein said identifying identifies the predetermined element through use of predetermined tags.
3. A method as recited in claim 2, wherein the markup language page is an HTML page, and wherein the predetermined tags are HTML tags.
4. A method as recited in claim 1, wherein the second network address
- 15 has a hostname pertaining to the intermediate server.
5. A method as recited in claim 4, wherein the second network address further has a suffix that includes at least a hostname of the first network address.
6. A method as recited in claim 4, wherein the second network address
- 20 further has a suffix that includes at least the first network address.
7. A method as recited in any of claims 1-6, wherein the predetermined element is an applet.
8. A method as recited in claim 1, wherein the predetermined element is a Java applet.
- 25 9. A method as recited in claim 1, wherein the predetermined element is a Java object.
10. A method for processing resource requests provided to an intermediate server from a client via a computer network, said method comprising the acts of:
 - 30 receiving a resource request for a particular resource, the resource request being provided to the intermediate server from the client via the computer network;
 - extracting a destination server from the resource request;
 - requesting the particular resource from the destination server;

- receiving the particular resource from the destination server;
modifying the particular resource to redirect internal resource requests
to the intermediate server;
sending the modified particular resource to the client;
- 5 receiving an applet code request for an applet identified within the
modified particular resource;
requesting applet code for the applet from a remote server via the
computer network;
receiving the applet code from the remote server in response to said
10 requesting of the applet code;
modifying the applet code to redirect its external communications
through the intermediate server; and
sending the modified applet code to the client.
11. A method as recited in claim 10, wherein the applet code is bytecode
15 for the applet.
12. A method as recited in claim 10, wherein said modifying of the
particular resource modifies at least one source file address within the
particular resource, the at least one source file address pertaining to the
applet.
- 20 13. A method as recited in claim 12, wherein the particular resource is a
markup language document.
14. A method as recited in claim 13, wherein the markup language
document is a HTML page.
15. A method as recited in claim 10, wherein the computer network
25 includes at least a portion of the Internet.
16. A method as recited in claim 10, wherein the remote server is the
destination server.
17. A method as recited in any of claims 10-16, wherein said modifying of
the particular resource comprises the acts of:
- 30 identifying, within the particular resource, a predetermined element that
includes at least a first network address; and
modifying the first network address within the predetermined element
of the particular resource to a second network address that pertains to the
intermediate server.

18. A method as recited in claim 17, wherein said identifying identifies the predetermined element through use of predetermined tags.
19. A method as recited in claim 17, wherein the particular resource request is a markup language document.
- 5 20. A method as recited in claim 18, wherein the markup language document is a HTML page.
21. A method as recited in claim 20, wherein the predetermined tags are HTML tags.
22. A method as recited in claim 16, wherein the second network address
10 includes at least a hostname pertaining to the intermediate server.
23. A method as recited in claim 22, wherein the second network address further has a suffix that includes at least the first network address.
24. A method as recited in any of claims 16-23, wherein the predetermined element pertains to an applet.
- 15 25. A method as recited in claim 24, wherein the predetermined element is a predetermined attribute of the applet.
26. A method for communicating between a client and a server, the client including an applet, said method comprising:
- determining whether a socket connection between the applet at the
20 client and server is available;
- establishing a socket connection between the applet at the client and the server when said determining determines that a socket connection is available;
- establishing a pair of unidirectional secure connections provided by
25 said communication layer when said determining determines that a socket connection is not available; and
- thereafter communicating data between the applet at the client and the server using whichever of the socket connection and the pair of unidirectional secure connections has been established.
- 30 27. A method as recited in claim 26, wherein the pair of unidirectional secure connections are URL connections.
28. A system for communicating between a client and a server, said system comprising:

a communication layer at a client, said communication layer transforming one or more socket connections into a pair of unidirectional secure URL connections;

an applet operating at the client to perform operations and to create at least one socket connection with said communication layer; and

a server operatively connected with the pair of unidirectional secure URL connections, said server communicating with said applet via the pair of unidirectional secure URL connections provided by said communication layer.

29. A system as recited in claim 28, wherein communications between said server and said applet are performed using packets of data.

30. A system as recited in claim 29, wherein the packets include a header and data, and wherein the header includes at least a browser identifier and a socket identifier.

31. A system as recited in claim 28, wherein said server includes one or more server socket connections with remote content servers, and wherein said server receives packets from said applet via one of the unidirectional secure URL connections provided by said communication layer, and wherein said server directs the packets received to an appropriate one of the server socket connections based on at least the browser identifier and the socket identifier.

32. A system as recited in claim 28, wherein communications between said server and said applet are performed using packets of data,

wherein the packets include a header and data, and the header includes at least a socket identifier,

wherein said server includes one or more server socket connections with remote content servers, and

wherein said server receives packets from said applet via one of the unidirectional secure URL connections provided by said communication layer, and directs the packets received to an appropriate one of the server socket connections based on at least the socket identifier.

33. A system as recited in claim 28, wherein said communication layer comprises:

an outgoing queue for buffering outgoing packets of data received from the at least one socket connection with said applet; and

a sender thread that sends the packets of data stored in said outgoing queue to said server.

- 5 34. A system as recited in claim 33, wherein said communication layer further comprises:

a reader thread that receives packets of data being sent from said server to said applet via said communication layer; and

- 10 at least one incoming queue for buffering incoming packets of data received from said server.

35. A system as recited in claim 34, wherein said at least one incoming queue operatively connects with the at least one socket connection of said applet.

- 15 36. A system as recited in claim 34, wherein flow control messages are sent between said communication layer and said server.

37. A system as recited in claim 34, wherein acknowledgement messages are sent between said communication layer and said server.

- 20 38. A system as recited in claim 28, wherein one of the pair of the unidirectional secure URL connections is an outgoing unidirectional secure URL connection, and the other of the pair of the unidirectional secure URL connections is an incoming unidirectional secure URL connection.

39. A system as recited in claim 38, wherein the outgoing unidirectional secure URL connection is a transient connection.

- 25 40. A system as recited in claim 39, wherein the incoming unidirectional secure URL connection is a persistent connection.

41. A system as recited in claim 28, wherein said communication layer is an applet communication layer.

42. A system for communicating between a client and a server, said system comprising:

- 30 a plurality of browser applications, at least a plurality of said browser applications utilizing at least one operating applet and a communication layer, said applet operating to perform operations and to create at least one socket connection with said communication layer, and said communication layer for

each of said browsers operates to transform said socket connections into a pair of unidirectional URL connections; and

5 a server operatively connected with the pair of unidirectional URL connections associated with said communication layer associated with each of said plurality of browser applications, said server communicating with said at least one operating applet of said plurality of browser applications via the pair of unidirectional URL connections provided by said communication layer corresponding thereto.

43. A system as recited in claim 42, wherein the pair of unidirectional URL
10 connections are secure connections.

44. A system as recited in claim 42, wherein the pair of unidirectional URL connections are HTTPS connections.

45. A system as recited in claim 42, wherein said server is an intermediate server that is provided as an intermediary between said browser application
15 and remote content servers.

46. A system as recited in claim 42, wherein one of the pair of unidirectional URL connections is a transient connection, and the other of the pair of unidirectional URL connections is a persistent connection.

47. A system for communicating between a client and a server, said
20 system comprising:

a plurality of browser applications, at least a plurality of said browser applications utilizing at least one operating applet and a communication layer, said applet operating to perform operations and to create at least one socket connection with said communication layer, and said communication layer for
25 each of said browsers operates to form an intermediate socket connection or to transform said socket connections into a pair of unidirectional URL connections; and

a server operatively connected with the intermediate socket connection or the pair of unidirectional URL connections associated with said
30 communication layer associated with each of said plurality of browser applications, said server communicating with said at least one operating applet of said plurality of browser applications via the intermediate socket connection or the pair of unidirectional URL connections provided by said communication layer corresponding thereto.

48. A system as recited in claim 47, wherein said server communicates with said at least one operating applet of said plurality of browser applications via the intermediate socket connection if such connection can be established, otherwise via the pair of unidirectional URL connections.

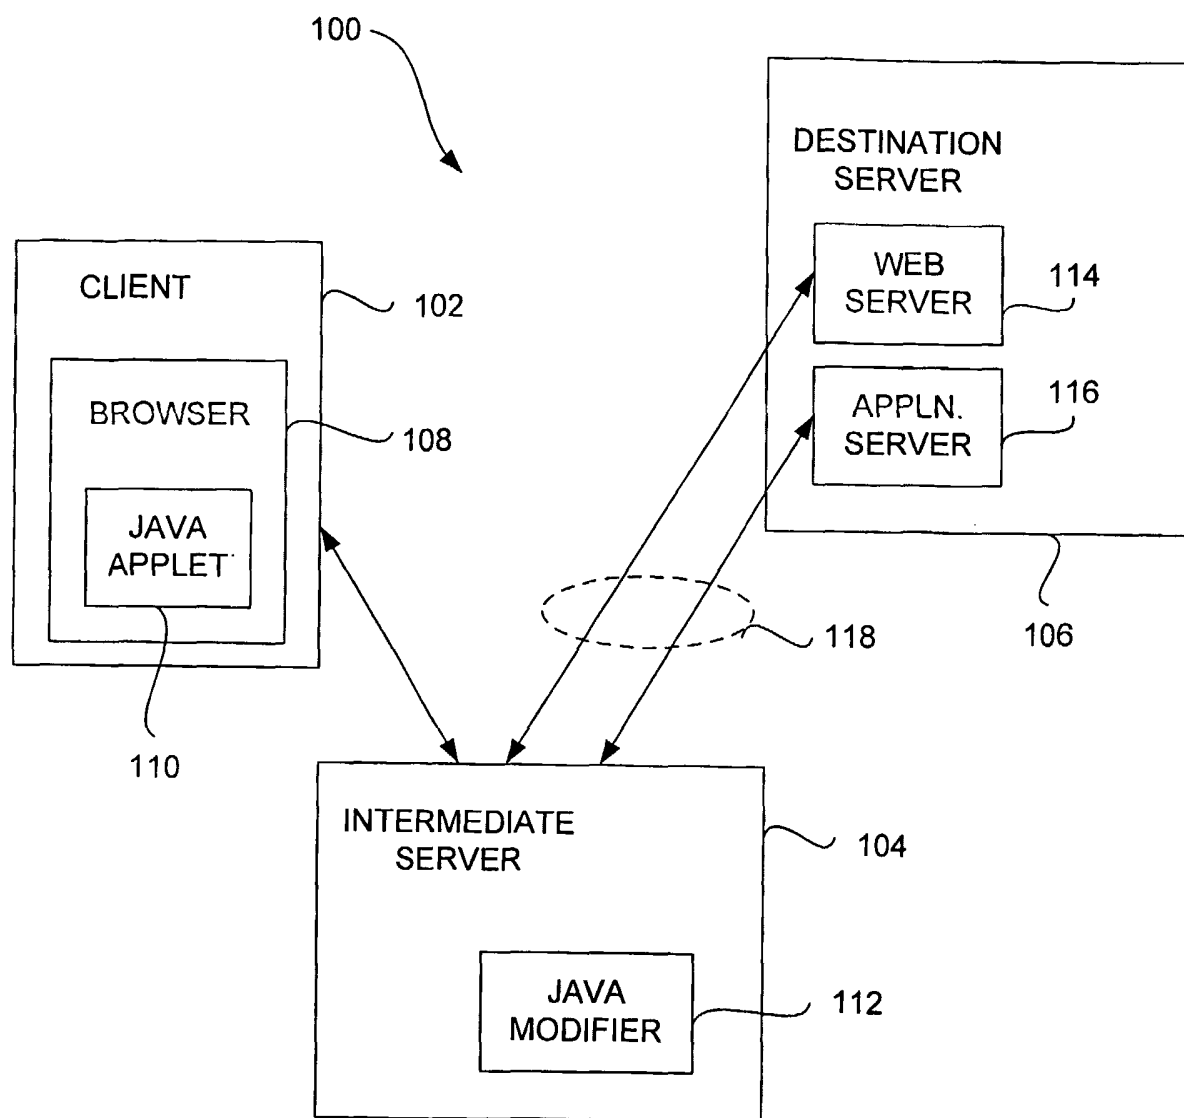
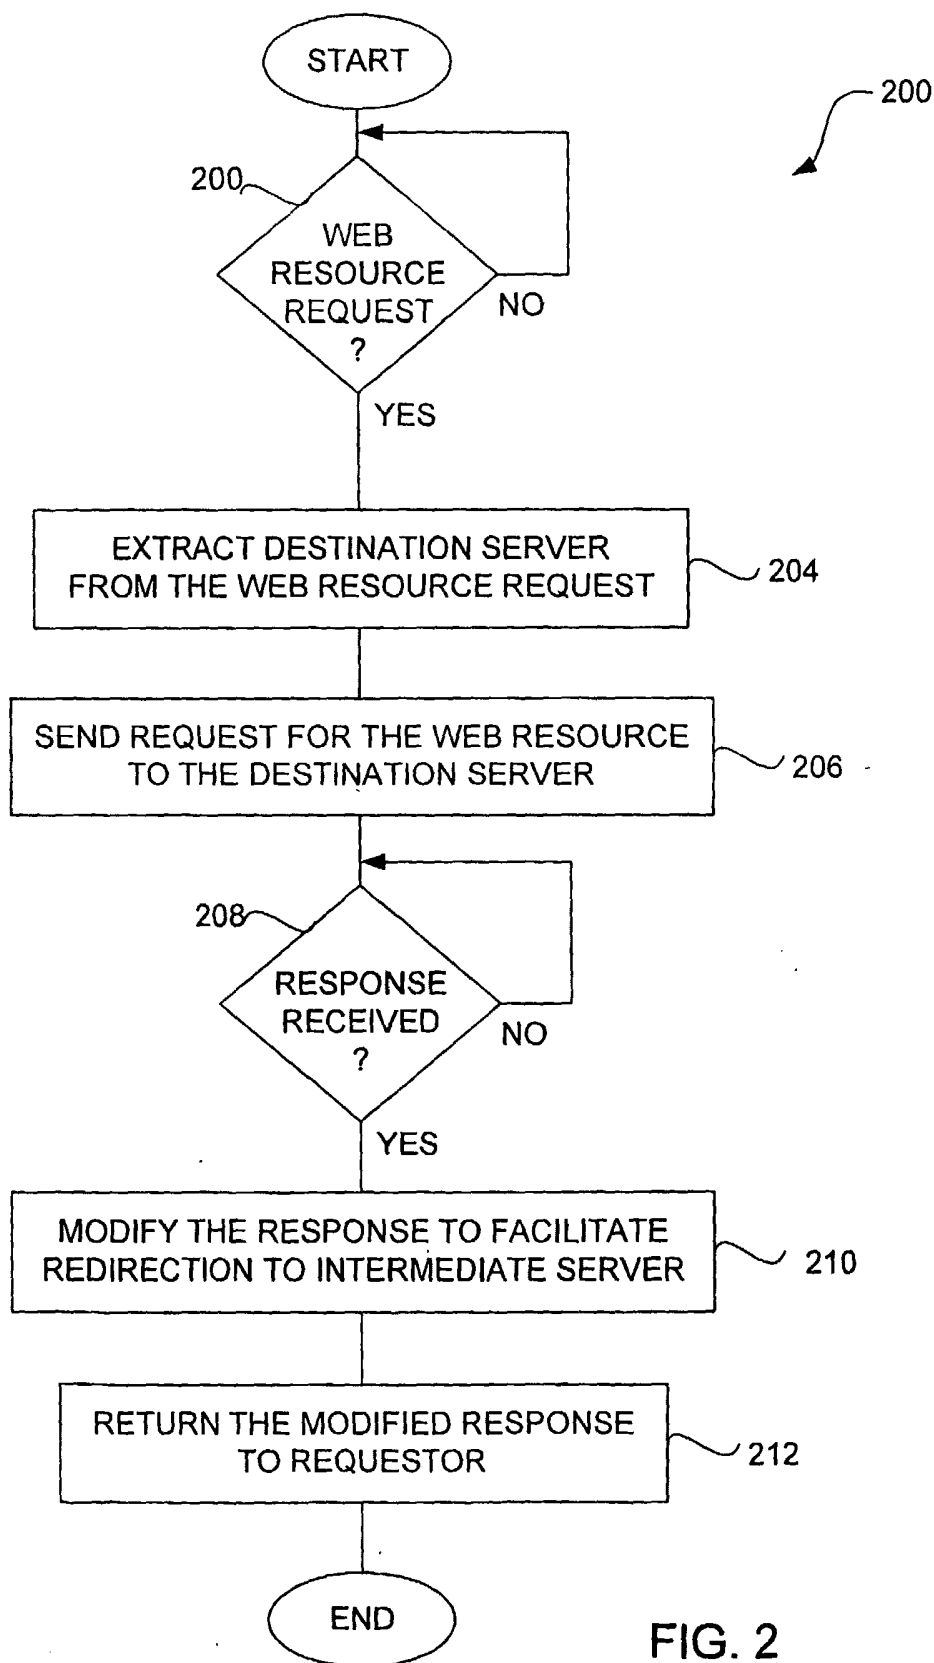


FIG. 1



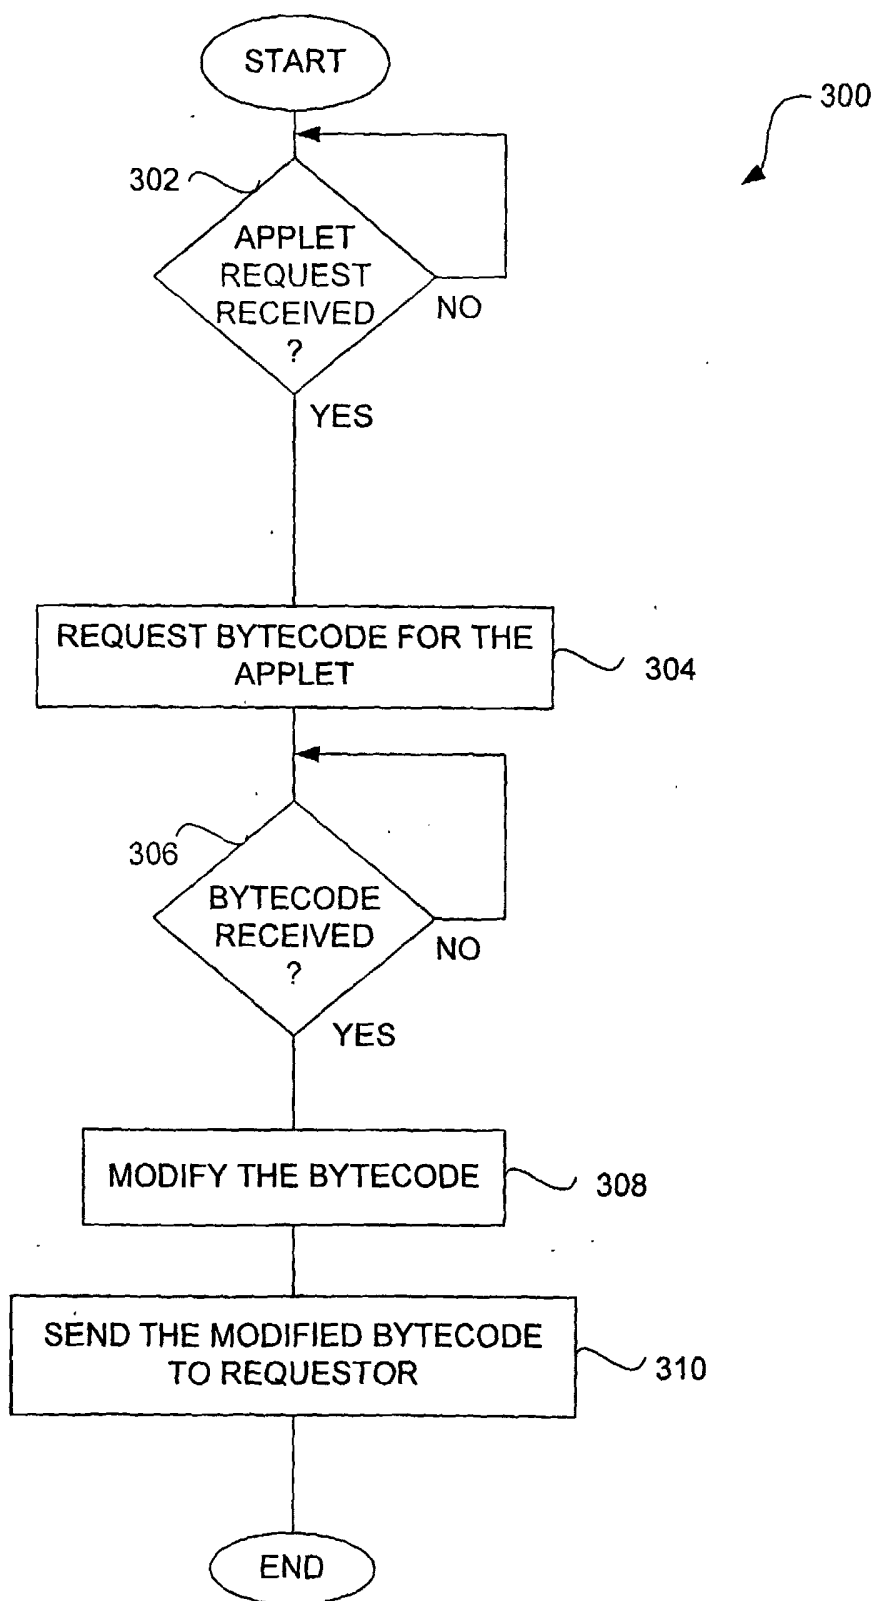
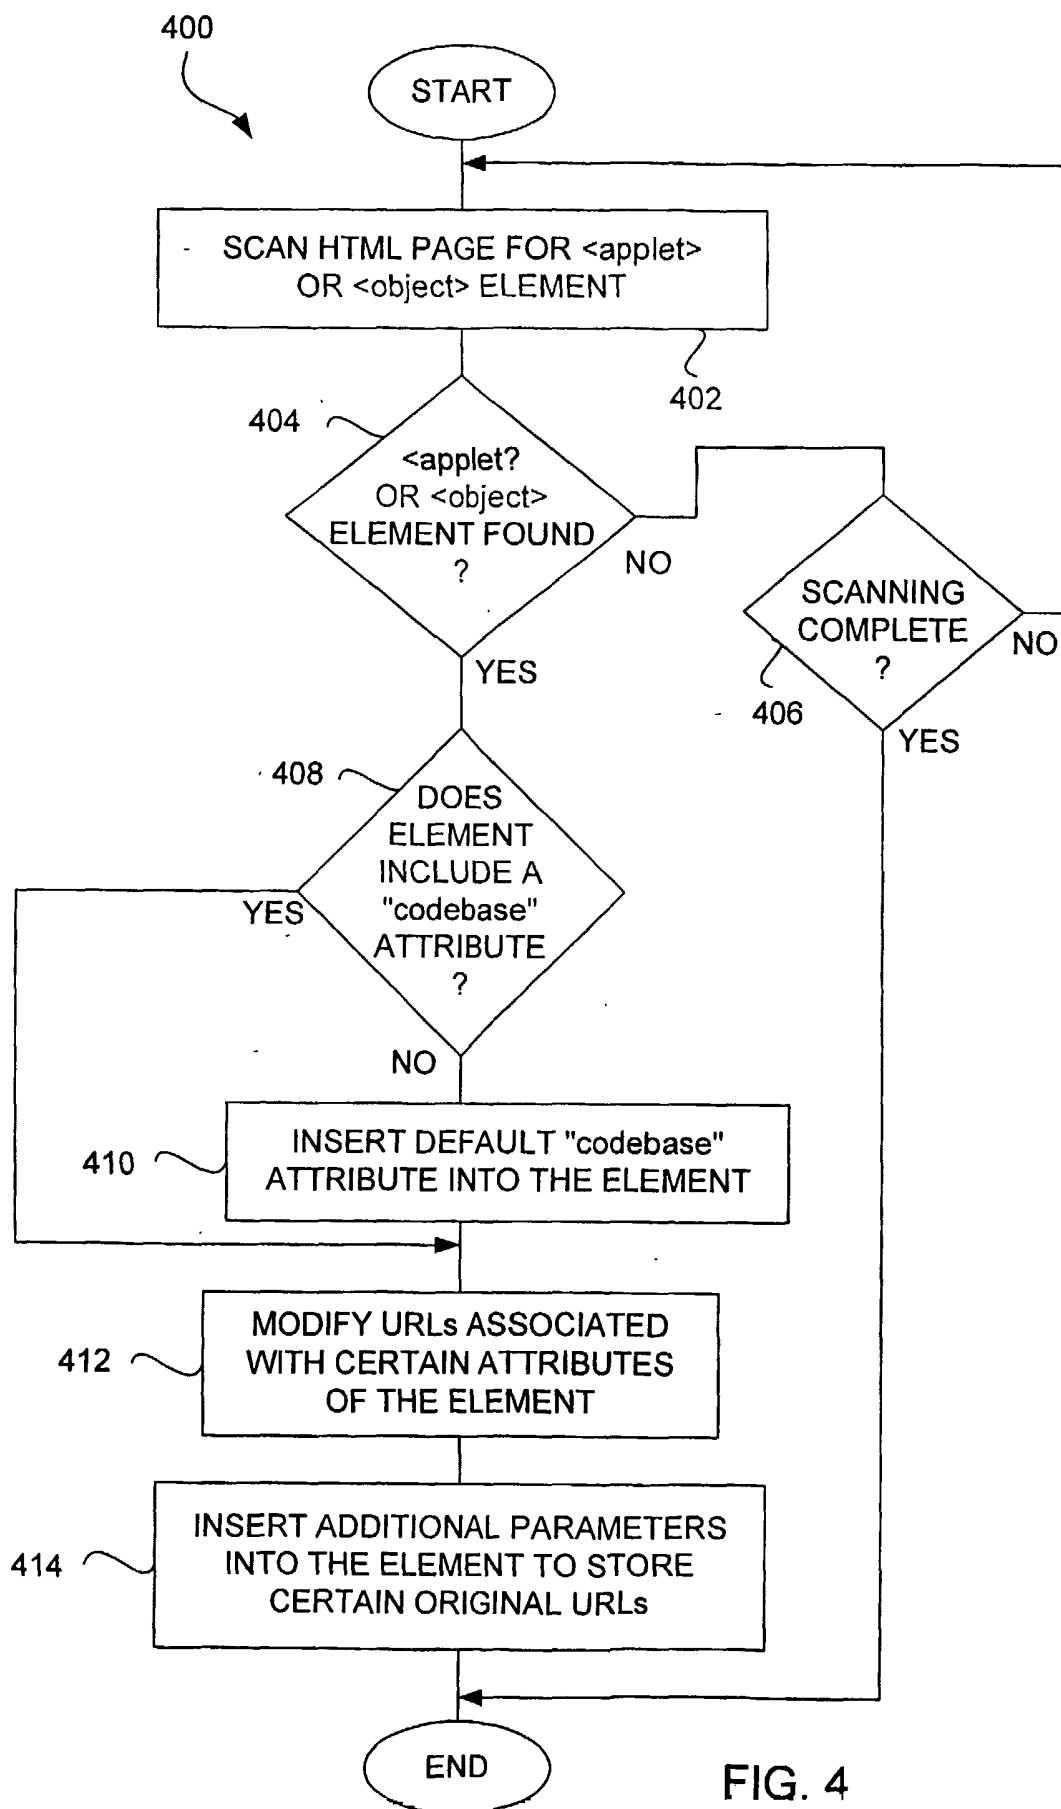


FIG. 3



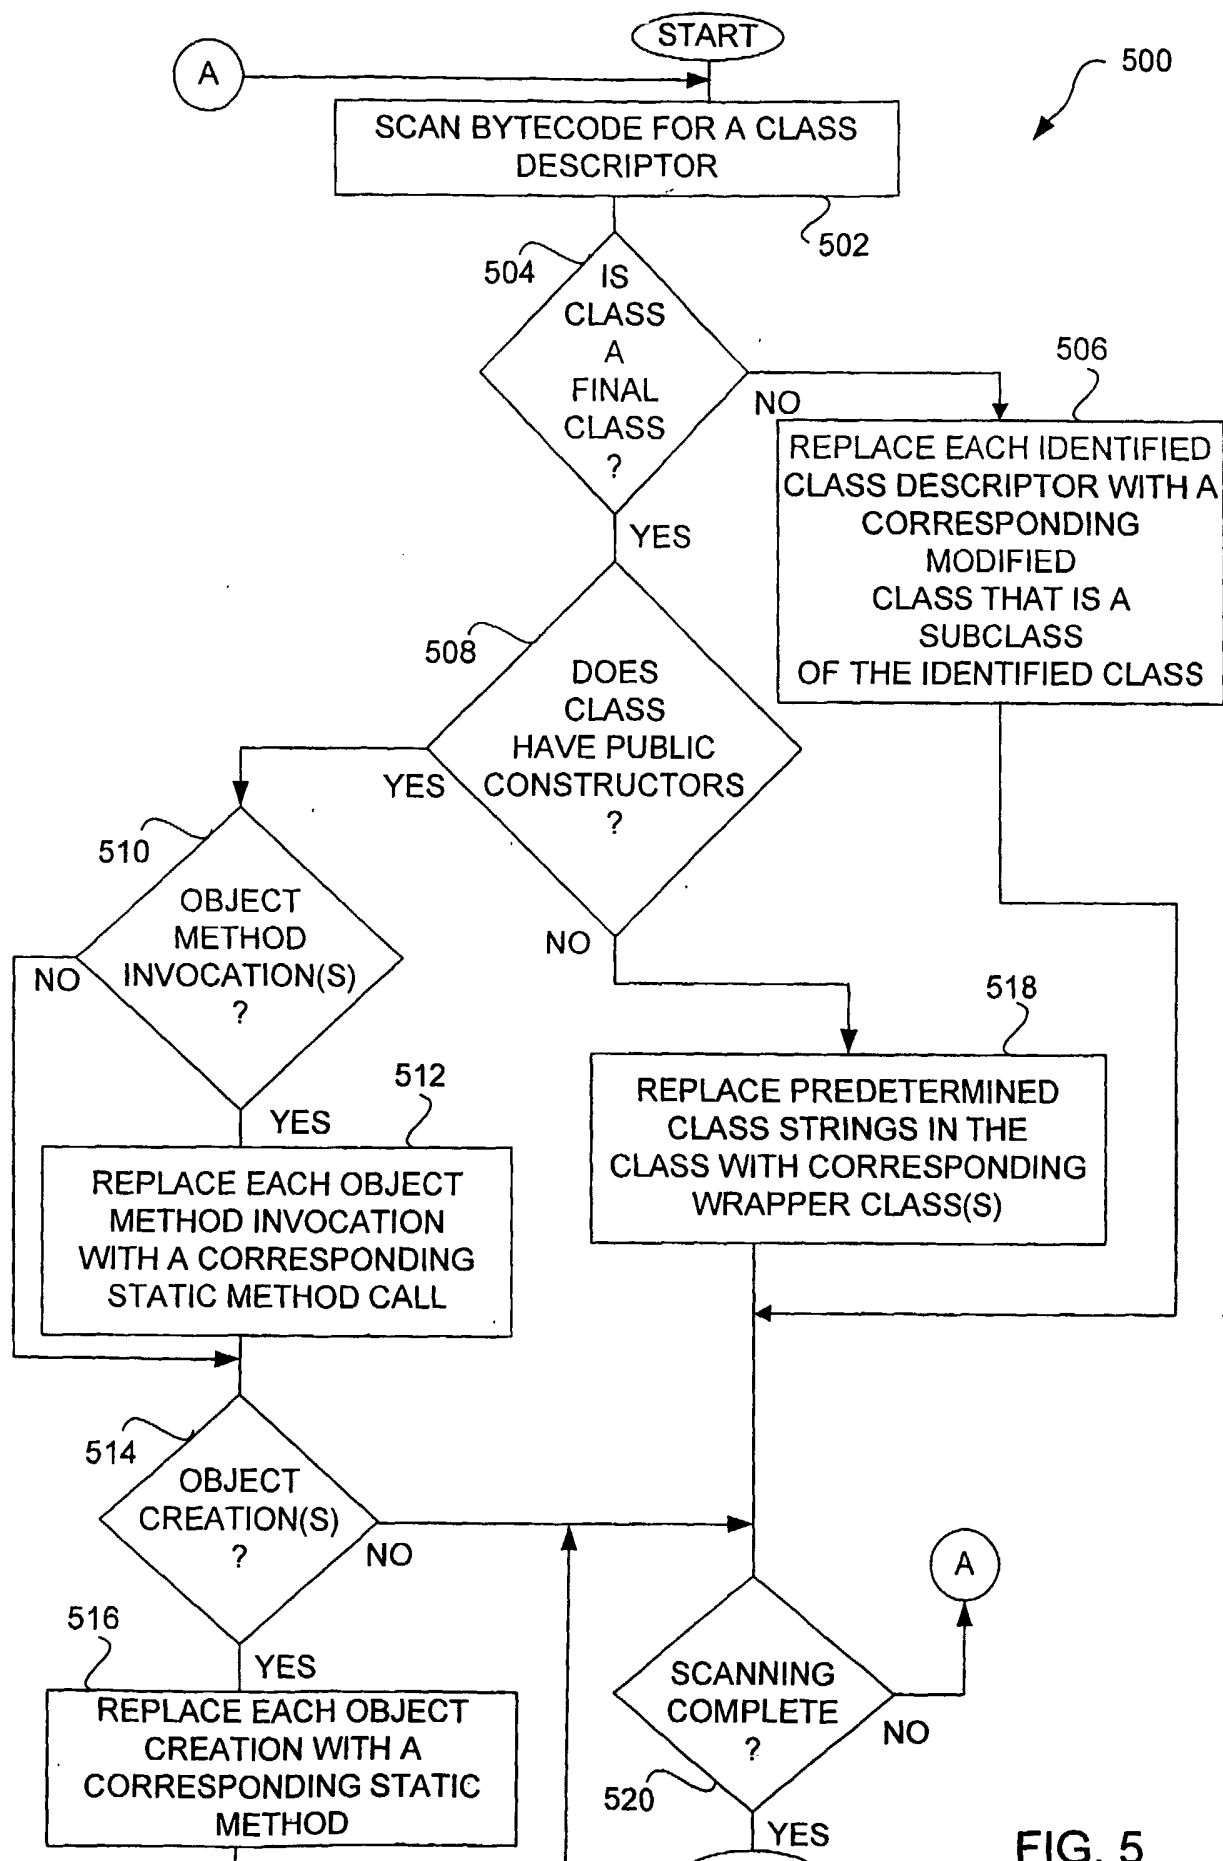


FIG. 5

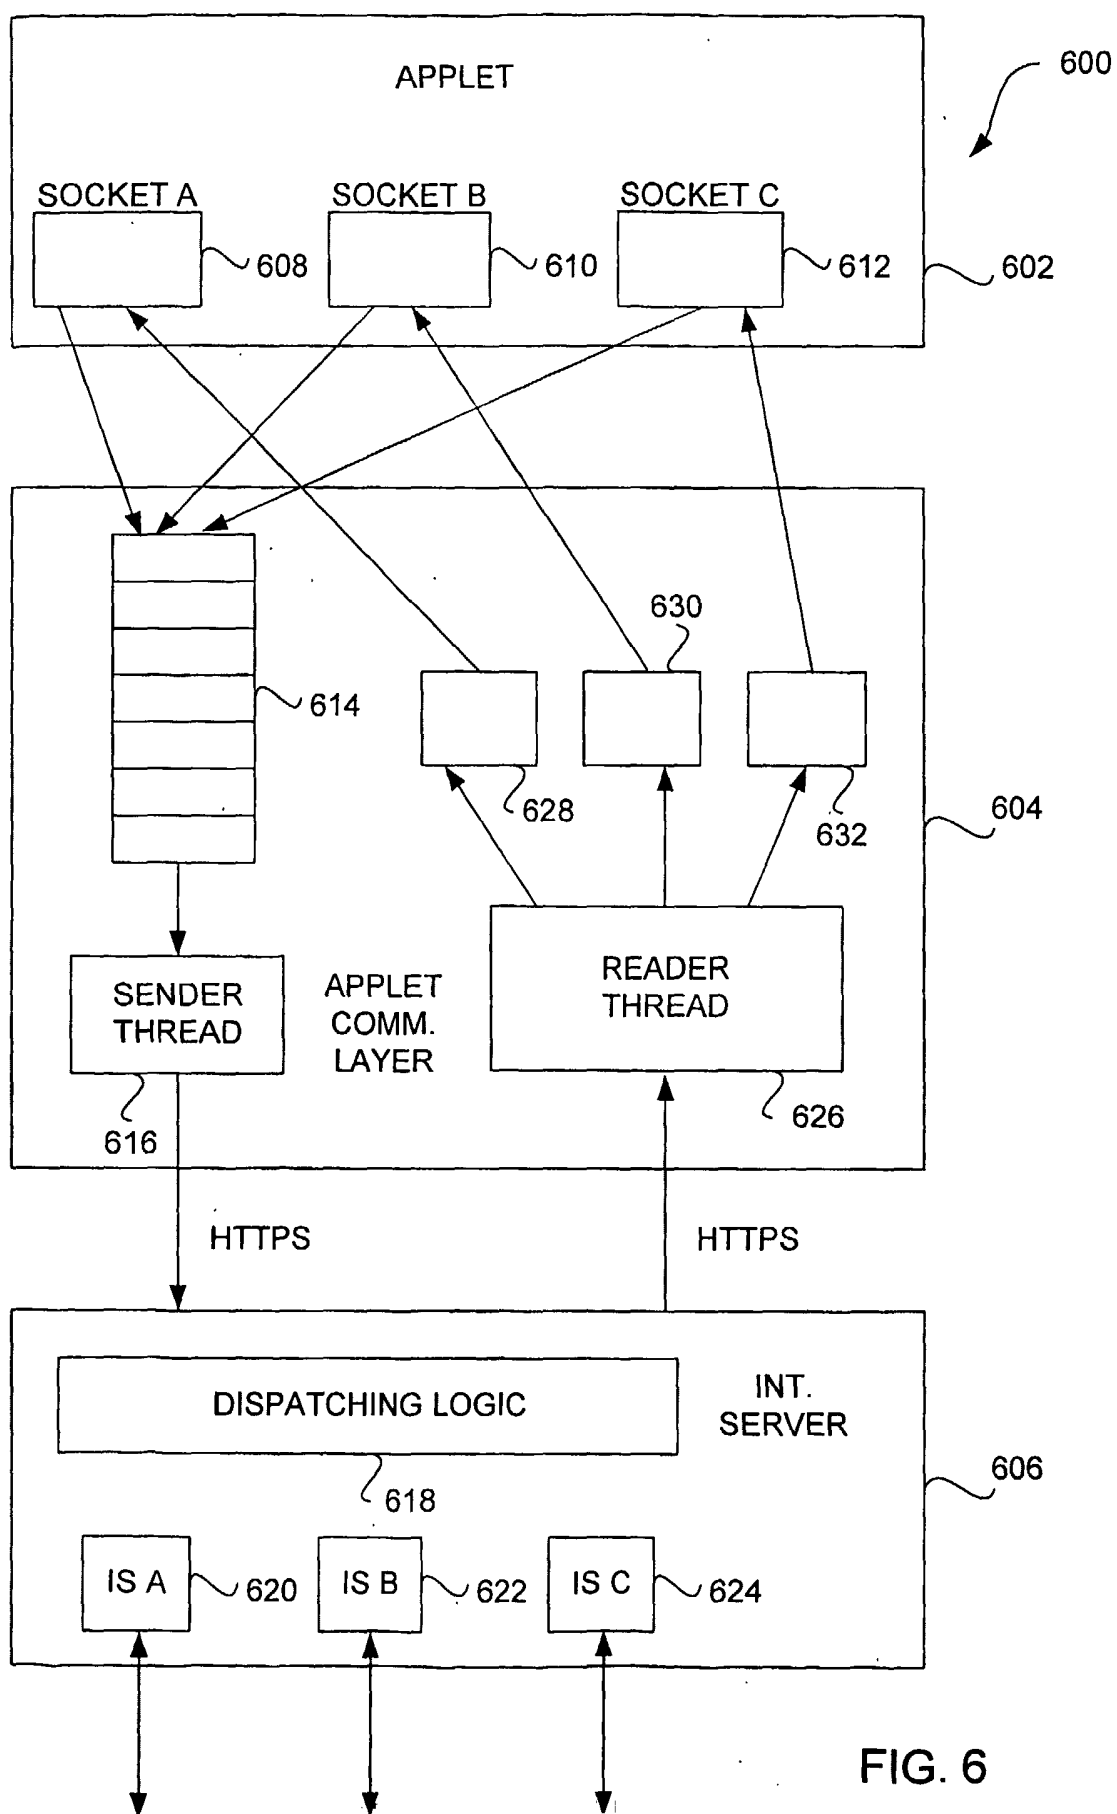


FIG. 6

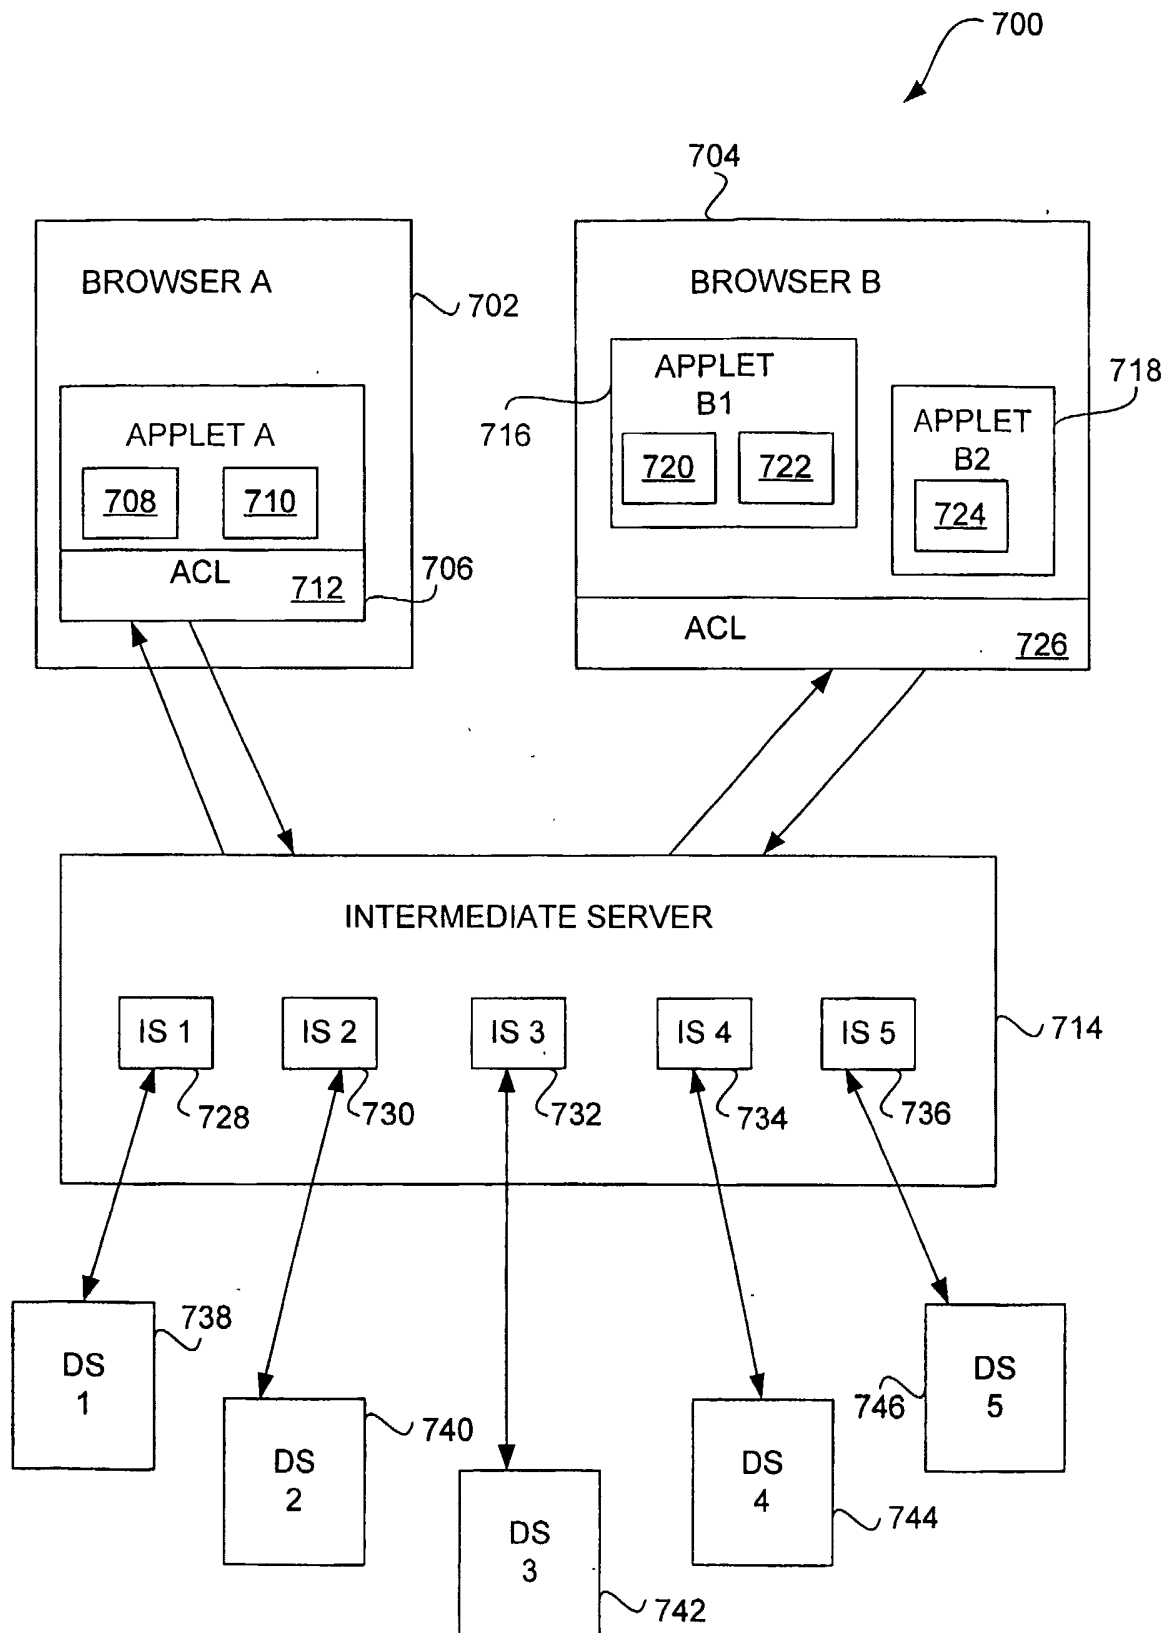


FIG. 7

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US03/17934

A. CLASSIFICATION OF SUBJECT MATTER

IPC(7) : G06F 15/16
US CL : 713/176

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. : 713/176; 709/203; 707/201,523; 705/71

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
Please See Continuation Sheet

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X --- Y	US 2002/0007393 A1 (HAMEL) 17 January 2002, abstract, Fig. 1, elements 115, 150 and 160, col. page 30, paragraph 30-36	1,10,26,28,42 and 47 ----- 2-6, 8-9, 11-16, 22-23, 27, 29-40, 43-46 and 48
Y	US 6,298,356 B1 (JAWAHAR et al) 02 October 2001, the entire document.	2-6, 8-9, 11-16, 22-23, 27, 29-41, 43-45 and 48
A	US 6,029,182 A (NEHAB et al) 22 February 2000, the entire document.	1, 10, 26, 28, 42 and 47

☐ 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 application or patent 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
"I"	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

22 October 2003 (22.10.2003)

Date of mailing of the international search report

13 NOV 2003

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US
Commissioner for Patents
P.O. Box 1450
Alexandria, Virginia 22313-1450

Facsimile No. (703)305-3230

Authorized officer

Ayaz Sheikh

Telephone No. 703-305-9648

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US03/17934

Box I Observations where certain claims were found unsearchable (Continuation of Item 1 of first sheet)

This international report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claim Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☐ Claim Nos.:
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
3. ☒ Claim Nos.: 7,17-21,24 and 25
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box II Observations where unity of invention is lacking (Continuation of Item 2 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fee, this Authority did not invite payment of any additional fee.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

☐
☐

The additional search fees were accompanied by the applicant's protest.

No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

PCT/US03/17934

Continuation of B. FIELDS SEARCHED Item 3:

User West, Dialog and Proquest NPL search. Search Terms; secure adj remote access same private adj network and applet or HTML
"in browser"