

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
10 July 2008 (10.07.2008)

PCT

(10) International Publication Number
WO 2008/080525 A1

(51) International Patent Classification:
G06F 11/30 (2006.01) **G06F 9/44** (2006.01)

(21) International Application Number:
PCT/EP2007/010882

(22) International Filing Date:
12 December 2007 (12.12.2007)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/647,979 29 December 2006 (29.12.2006) US

(71) Applicant (for all designated States except US): **SAP AG**
[DE/DE]; Dietmar-Hopp-Allee 16, 69190 Walldorf (DE).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **BONEV, Pavel**
[BG/BG]; j.k. Svoboda bl.4, ent.B, app.30, 1231 Sofia
(BG). **STANEV, Georgi** [BG/BG]; Kozylak Str. no.21,
ap.1, 1407 Sofia (BG). **DROSHEV, Mladen** [BG/BG];
"Bulkston", bl.6, ent.4, ap.49, 1618 Sofia (BG).

(74) Agent: **SCHIUMA, Daniele**; Müller-Boré & Partner,
Grafinger Str. 2, 81671 München (DE).

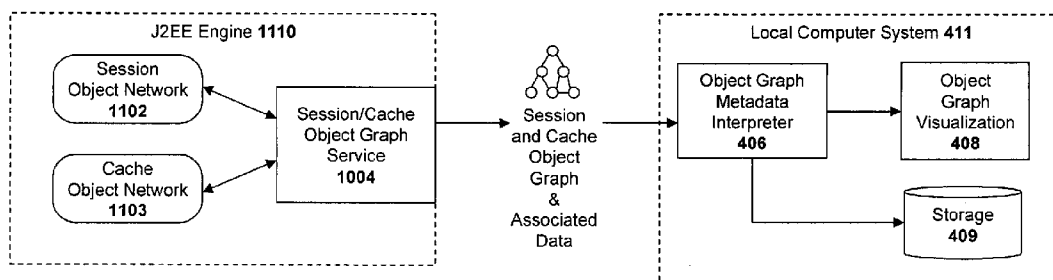
(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (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

(54) Title: GRAPHICAL USER INTERFACE SYSTEM AND METHOD FOR PRESENTING INFORMATION RELATED TO SESSION AND CACHE OBJECTS



(57) Abstract: A system and method for visualizing information related to session and cache objects within an object network. For example, a computer-implemented method according to one embodiment comprises: analyzing relationships between session and/or cache objects within a network of sessions and caches to determine an object network structure for the session and/or cache objects; generating object graph data representing the object network structure; serializing the object graph data and transmitting the object graph data over a network to a requesting computer; and interpreting the object graph data to render a view of the session and/or cache object network structure in a graphical user interface.

WO 2008/080525 A1

GRAPHICAL USER INTERFACE SYSTEM
AND METHOD FOR PRESENTING INFORMATION
RELATED TO SESSION AND CACHE OBJECTS

BACKGROUND

Field of the Invention

This invention relates generally to the field of data processing systems. More particularly, the invention relates to a graphical user interface system and method for presenting information related to session and cache object graphs.

Description of the Related Art

Multi-Tiered Enterprise Computing Systems

Traditional client-server systems employed a two-tiered architecture such as that illustrated in **Figure 1a**. Applications 102 executed on the client side 100 of the two-tiered architecture are comprised of a monolithic set of program code including a graphical user interface component, presentation logic, business logic and a network interface that enables the client 100 to communicate over a network 103 with one or more servers 101. A database 104 maintained on the server 101 provides non-volatile or “persistent” storage for the data accessed and/or processed by the application 102.

The “business logic” component of the application represents the core program code of the application, i.e., the rules governing the underlying business process (or other functionality) provided by the application. The “presentation logic” describes the specific manner in which the results of the business logic are

formatted for display on the user interface. The “database” 104 includes data access logic used by the business logic to store and retrieve data.

The limitations of the two-tiered architecture illustrated in **Figure 1a** become apparent when employed within a large enterprise. For example, installing and maintaining up-to-date client-side applications on a large number of different clients is a difficult task, even with the aid of automated administration tools. Moreover, a tight coupling of business logic, presentation logic and the user interface logic makes the client-side code very brittle. Changing the client-side user interface of such applications is extremely hard without breaking the business logic, and vice versa. This problem is aggravated by the fact that, in a dynamic enterprise environment, the business logic may be changed frequently in response to changing business rules. Accordingly, the two-tiered architecture is an inefficient solution for enterprise systems.

In response to limitations associated with the two-tiered client-server architecture, a multi-tiered architecture has been developed, as illustrated in **Figure 1b**. In the multi-tiered system, the presentation logic 121, business logic 122 and database 123 are logically separated from the user interface 120 of the application. These layers are moved off of the client 125 to one or more dedicated servers on the network 103. For example, the presentation logic 121, the business logic 122, and the database 123 may each be maintained on separate servers, 126, 127 and 128, respectively.

This separation of logical components and the user interface provides a more flexible and scalable architecture compared to that provided by the two-tier model. For example, the separation ensures that all clients 125 share a single implementation of business logic 122. If business rules change, changing the current implementation of business logic 122 to a new version may not require updating any client-side program code. In addition, presentation logic 121 may be provided which generates code for a variety of different user interfaces 120, which may be standard browsers such as Internet Explorer® or Netscape Navigator®.

The multi-tiered architecture illustrated in **Figure 1b** may be implemented using a variety of different application technologies at each of the layers of the multi-tier architecture, including those based on the Java 2 Enterprise Edition™ (“J2EE”) standard, the Microsoft .NET standard and/or the Advanced Business Application Programming (“ABAP”) standard developed by SAP AG. For example, as described below, in a J2EE environment, the business layer 122, which handles the core business logic of the application, is comprised of Enterprise Java Bean (“EJB”) components with support for EJB containers. Within a J2EE environment, the presentation layer 121 is responsible for generating servlets and Java Server Pages (“JSP”) interpretable by different types of browsers at the user interface layer 120.

J2EE Application Server Architecture

Figure 2 illustrates a typical J2EE application server 200 in which the presentation layer is implemented by a “Web container” 211 and the business layer is implemented by an Enterprise Java Bean (“EJB”) container 201. Containers are runtime environments which provide standard common services 219, 209 to runtime components. For example, the Java Naming and Directory Interface (“JNDI”) is a service that provides application components with methods for performing standard naming and directory services. Containers also provide unified access to enterprise information systems 217 such as relational databases through the Java Database Connectivity (“JDBC”) service, and legacy computer systems through the J2EE Connector Architecture (“JCA”) service. In addition, containers provide a declarative mechanism for configuring application components at deployment time through the use of deployment descriptors.

As illustrated in **Figure 2**, each layer of the J2EE architecture includes multiple containers. The Web container 211, for example, is itself comprised of a servlet container 215 for processing servlets and a Java Server Pages (“JSP”) container 216 for processing Java server pages. The EJB container 201 includes three different containers for supporting three different types of enterprise Java beans: a session bean container 205 for session beans, a entity bean container 206 for entity beans, and a message driven bean container 207 for message driven beans. A more detailed description of J2EE containers and J2EE services can be found in RAGAE GHALY AND KRISHNA KOTHAPALLI, SAMS TEACH YOURSELF EJB IN 21 DAYS (2003) (see, e.g., pages 353-376).

Object-Oriented Computer Systems

The computer systems described above consist of many smaller pieces of program code referred to as “objects” which interact each other. For example in a computer program for booking cars at least three objects are required for storing the relevant information: one for the person who makes the booking (name, credit card number etc), one for the booked car (model, engine, class, etc) and another for the booking itself (booking date, return date, etc).

In most cases, objects reference other objects to form very complex object networks. Sometimes information about the structure of an object network is needed without knowledge about the real data and semantics of the objects in the network. By way of analogy, within a genealogic tree, to determine structural data such as how many people are linked, how clustered the tree is, etc, it is not necessary to know the details about each person – just the properties of the network itself. For that purpose, a consistent, flexible way to represent the object network structure would be desirable.

SUMMARY

A system, method and computer program product (particularly embodied in a machine-readable medium, as a signal and/or as a data stream) for visualizing information related to session and cache objects within an object network. For example, a computer-implemented method according to one embodiment comprises: analyzing relationships between session and/or cache objects within a network of sessions and caches to determine an object network structure for the session and/or cache objects; generating object graph data representing the object network structure; serializing the object graph data and transmitting the object graph data over a network to a requesting computer; and interpreting the object graph data to render a view of the session and/or cache object network structure in a graphical user interface. Accordingly, man-machine interaction can be improved, particularly as the graphical user interface allows a user to more readily understand or interpret the session and/or cache object network structure.

BRIEF DESCRIPTION OF THE DRAWINGS

A better understanding of the present invention can be obtained from the following detailed description in conjunction with the following drawings, in which:

FIG. 1a illustrates a traditional two-tier client-server architecture.

FIG. 1b illustrates a prior art multi-tier client-server architecture.

FIG. 2 illustrates a multi-tiered application server architecture according to the Java 2 Enterprise Edition ("J2EE") standard.

FIG. 3 illustrates an exemplary object graph structure.

FIG. 4 illustrates an architecture for processing object graph data according to one embodiment of the invention.

FIG. 5 illustrates a graphical user interface according to one embodiment of the invention.

FIG. 6 illustrates session data within a prior art J2EE application server architecture.

FIG. 7 illustrates a hierarchical session domain architecture according to one embodiment of the invention.

FIG. 8 illustrates an exemplary set of session objects.

FIG. 9a illustrates hierarchical session domains for the session object in FIG. 8.

FIG. 9b illustrates session domains for an HTTP session context.

FIG. 10 illustrates data contained within one embodiment of a session domain.

FIG. 11 illustrates a session/cache object graph service employed in one embodiment of the invention.

FIGS. 12a-b illustrate a method for measuring differences in memory consumption between object.

FIG. 13 illustrates an application server architecture according to one embodiment of the invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Described below is a system and method for processing object graphs. Throughout the description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, to one skilled in the art that the present invention may be practiced without some of these specific details. In other instances, well-known structures and devices are shown in block diagram form to avoid obscuring the underlying principles of the present invention.

A. System and Method for Processing Object Graphs

Figure 3 illustrates an exemplary network of objects 300-305 which will be used to describe the various embodiments of the invention. Each of the objects may contain different types of data and methods. In the illustrated example, object 300 references objects 301 and 302. For example, object 300 points to objects 301 and 302 (e.g., object 300 may include object 301, 302 instances as its properties). The contained object state is part of the state of the object – container. Similarly, object 302 references objects 301 and 304-305 and object 301 references objects 303-304.

As mentioned above, information about the structure of an object network such as the one shown in **Figure 3** may be needed without the need for the data and semantics of the objects themselves. For example, information related to how each of the objects reference one another may be needed rather than the

underlying data contained within the objects. In these situations, it is inefficient to transfer all of the data when all that is required is information related to the relationships between the objects.

To address these needs, one embodiment of the invention analyzes object oriented program code and generates a math graph structure to represent the object network. The math graph structure may then be viewed within a visualization tool and/or stored within a mass storage device (as described below). The graph structure consists of one or many "nodes," each of which represents an object, and zero or many "arcs" (such as arcs 310) between the nodes. In an embodiment which uses Java program code (or a similar type of program code), an oriented graph is used because Java object "has a" relationships are parent-child relationships (e.g., a car "has a" make, model, year, and engine). In this embodiment, every object in computer memory is represented by one graph node and every object-to-object relationship is represented by one arc.

Figure 4 illustrates one embodiment of an architecture for performing the foregoing operations. A local computer system 411 requests object graph information related to object-oriented program code 402 executed on a remote computer system 410. An object graph processing service 404 extracts information from the object-oriented program code 402 to generate the object graph and other types of data described herein. The object graph and other data are then transmitted over a network to the requesting computer system 411.

Object graph interpreter logic 406 on the requesting computer system 411 interprets the serialized object graph and other data to graphically display the results within an object graph graphical user interface 408 sometimes referred to herein as a “visualization tool” (such as the one described below). In addition, the object graph and other data may be stored within a storage medium 409 such as a file system or database for later retrieval and viewing.

In one embodiment, the object graph processing logic 404 generates and/or collects additional information about objects (nodes) and references (arcs) and transmits the additional information with the object graph information. For example, in one embodiment, the following additional information is generated and/or collected:

1. The “memory size” of each object. This is the relative amount of random access memory consumed by each object. In one embodiment, the object graph processing logic 404 calculates the memory size in terms of a weighted value rather than an absolute value. In one particular implementation, a different weighted value is assigned to each of the generic types within the object (e.g., based on the relative memory consumption of those types). For example, an integer (INT) may be assigned a weight of 4 whereas a Boolean type may be assigned a weight of 1. Various other weights may be assigned to other generic types including, for example, char, float, double, byte, short, long and string. In one embodiment, the following weights are assigned to each of the generic types: boolean = 1 byte; byte = 1 byte; char = 2 bytes; short = 2 bytes; int = 4 bytes; long = 8 bytes; float = 4 bytes; double = 8 bytes.

2. The name of the Java type (i.e., the class).
3. The object's unique identification code for the graph structure. In one embodiment, the object graph processing logic 404 assigns each object an integer value which uniquely identifies the object within the graph structure.
4. An indication as to whether the object is serializable in Java terms. As it is known in the art, a "serializable" object may be converted to a bit stream and transmitted over a network.
5. An indication as to whether the object is shareable. Certain computer systems designed by the assignee of the present application allow objects to be shared among multiple Java virtual machines (as described below). Thus, objects in these systems may be designated "shareable" or "non-shareable."

In addition, the following information is generated and/or collected for each of the references:

1. The reference name (i.e., the class field name).
2. An indication as to whether the reference is a transient field in Java terms. As is known in the art, a "transient" field is one which is neither serialized or persisted.
3. The identity of two nodes initiating this relationship (i.e., the parent and child nodes). In one embodiment, this is accomplished using the objects' unique identification codes mentioned above.

The graph structure and additional information related to the object network are then serialized and transmitted as metadata to a requesting client.

This is advantageous because the graph structure and additional information can be observed as any other graph object using various different types of graph visualization tools. It can also be transmitted over a network such as the Internet (without sending the actual objects and data). It may also be saved on disk in a specified file type (e.g., a text/XML file) or as a Java serialized object file for later observation.

The object graph processing logic 404 may format and store the object graph data in a variety of ways while still complying with the underlying principles of the invention. For example, both list structures and matrix structures may be used, or a combination of both. List structures are often used for sparse graphs as they have smaller memory requirements whereas matrix structures provide faster access but may consume significant amounts of memory if the graph is very large.

List structure types may include an incidence list structure or an adjacency list structure. In an incidence list structure, the edges are represented by an array containing pairs (ordered if directed) of nodes (that the edge connects) and eventually weight and other data. In an adjacency list structure, each node (also sometimes referred to as a "vertex") has a list of which vertices it is adjacent to. This causes redundancy in an undirected graph: for example, if vertices A and B are adjacent, A's adjacency list contains B, while B's list contains A. Adjacency queries are faster, at the cost of extra storage space.

Matrix structures may include an incidence matrix structure in which the graph is represented by a matrix of E (edges) by V (vertices), where [edge, vertex] contains the edge's data and/or an adjacency matrix which is an N by N matrix, where N is the number of vertices in the graph. If there is an edge from some vertex X to some vertex Y, then the element $M_{x,y}$ is 1; otherwise it is 0. This type of matrix makes it easier to locate subgraphs, and to reverse graphs if needed. Other possible matrix types include a Laplacian matrix, Kirchhoff matrix, admittance matrix and a distance matrix.

It should be noted that the underlying principles of the invention are not limited to any particular format for the object graph data or to the specific architecture shown in **Figure 4**. For example, the “remote” system and “local” system do not necessarily need to be located in separate locations or on separate machines.

B. A Graphical User Interface for Object Graphs

After retrieving the object graph and other data associated with objects on the remote computer system, one embodiment of the invention employs a unique graphical visualization tool to observe the object graph and data. One embodiment of a visualization tool 408 is illustrated in **Figure 5**. In this embodiment, in response to user-selection of a particular node 520 (e.g., via a mouse or other cursor control device), the graph is drawn as a tree using the selected node 520 as the root. In one embodiment, the visualization tool 408 renders each of the child nodes further up in the hierarchy before their children

(i.e., it employs a breadth-first scan). However, the underlying principles of the invention are not limited to the order in which the nodes are rendered. If a node is referenced more than once the visualization tool 408 renders it each time as a child node but does not traverse its sub-tree more than once. One reason for this is that it avoids cyclic dependencies and potentially endless drawing.

As illustrated in **Figure 5**, in one embodiment of the visualization tool 408 employs different shapes, colors and graphics to indicate different types of nodes and arcs. For example, in one embodiment, the visualization tool uses a different filling color for nodes referenced once or many times. In **Figure 5**, an orange filling color is used to identify nodes referenced more than once, as indicated by entry 503 in the legend window 500, and a yellow filling color is used to identify nodes referenced only once.

In addition, as indicated in the legend window 500, a green filling color is used to identify “dummy” nodes 507. Dummy nodes are used to improve system performance when working with extremely large graphs. For example, a particular object graph may be very large (e.g., over 1 Gbyte) and may include hundreds of thousands of nodes. In this case, it may take a significant amount of time to traverse and render the entire graph. As such, in one embodiment, when a certain number of nodes have been traversed (e.g., 10,000), the object graph processing logic 404 and/or the visualization tool 408 may stop rendering and insert dummy nodes at the lowest level of the graph structure. The

remainder of the graph structure may be drawn upon user request (e.g., by selecting one or more of the dummy nodes).

In one embodiment, different shapes are used to distinguish between nodes with and without children. In the specific example shown in **Figure 5**, nodes without children are circular, as indicated by legend entry 501, and nodes with children are square with curved corners, as indicated by legend entry 502.

In addition, in the illustrated embodiment, different shape contour colors are used to mark nodes having different properties. For example, a light red contour color is used for nodes which are not shareable, as indicated by legend entry 504, and a dark red contour color is used for nodes which have children which are non-shareable, as indicated by legend entry 505. Various different contour colors may be used to identify different node properties (e.g., serializable nodes may be assigned another contour color). Of course, the underlying principles of the invention are not limited to any particular shapes and/or colors.

In addition, in the embodiment shown in **Figure 5**, the visualization tool 408 uses a graphic indicating multiple overlaid nodes for “compound” nodes. As used herein, a compound node is an array of nodes of the same type (class).

In one embodiment, the visualization tool 408 appends a label on each node. In the example shown in **Figure 5**, the label comprises the identification code assigned to each node. However, various other and/or additional

information may be displayed including, for example, the node weight and/or the percentage of the node's parent weight.

In one embodiment, the visualization tool 408 generates the tree dynamically, in response to user input. For example, the visualization tool 408 may expand/collapse the tree upon selection of a node. In addition, in one embodiment, the visualization tool 408 skips certain nodes which do not meet a user-specified criteria (e.g., displaying only nodes which are not shareable). In one embodiment, the visualization tool 408 skips nodes from a given type (class) (class fields with a given name, etc).

In one embodiment, the visualization tool 408 displays the additional data collected for each node in response to user input. For example, in **Figure 5**, a node information window 510 is generated when the user moves a cursor over a particular node. The node information window 510 may include any of the information described herein including, for example, the node ID, the type, the relative weight (i.e., "Native Size"), and whether the node is shareable.

As mentioned above, in one embodiment, additional metadata related to each of the object references is generated and/or collected such as the reference name (i.e., the class field name); an indication as to whether the reference is a transient field in Java terms; and the identity of two nodes initiating this relationship. This information may then be visually displayed within the visualization tool. For example, as indicated in **Figure 5**, lines representing

transient references may be provided in a different color than those for standard references. Various other other graphical features may be employed to convey information about object relationships (e.g., such as a dotted line for unexplored node references).

C. Session and Cache Object Networks

The assignee of the present application has developed advanced, hierarchical architectures for managing session objects and cache objects. See, e.g., Session Management Within a Multi-Tiered Enterprise Network, Serial No. ____, Filed December 28, 2004, for session object management and Distributed Cache Architecture, Serial No. 11/025,714, Filed December 28, 2005, for cache management. Each of these co-pending patent applications is assigned to the assignee of the present application and is incorporated herein by reference.

The following discussion will focus on embodiments related to session objects and data but the same general principles apply equally to cache objects and data. An application server manages session data associated with each client interaction. The session data defines the current conversational state between the client and the application server. For example, if the user is purchasing books from an online bookstore, the session data may define the current state of the user's "shopping cart" (e.g., the books the user has selected for purchase). Similarly, if the user is logged in to a particular Web portal (e.g.,

“My Yahoo”), the session data may define the state of the Web pages opened in the user’s Web browser.

In one embodiment, different types of session data are maintained at each logical layer of the application server. For example, referring to **Figure 6**, in response to a single client 630 request from a Web browser, an HTTP session object 641 (e.g., an instance of `javax.servlet.HttpSession`) containing HTTP session data is managed within the Web container 640 and (depending on the type of application) multiple session beans 651-653 may be managed within the EJB container 650. Additional HTTP session objects 642 and session beans 654-655 may be managed concurrently to define the conversational state with additional clients 631 for the same application (e.g., the shopping cart application) or other applications (e.g., a “search” application). Thus, the state of any given set of applications such as the user’s shopping cart may be spread across multiple containers.

One embodiment of the invention employs session management logic employs a hierarchical structure for storing different types of related session data within each enterprise application (e.g., HTTP sessions, session EJBs, etc). In particular, as illustrated in **Figure 7**, one embodiment of the session management layer groups related session objects within a logical hierarchy. At the top of the hierarchy is a session context object 700. “Session domains” are the base configurable objects within the session hierarchy and are positioned beneath the session context 700. The session domains represent abstract

storage units for sessions that have similar characteristics (e.g., similar life cycles, applications, etc).

In the example shown in **Figure 7**, session objects 711 and 712 are managed within session domain 710, and session objects 721 and 722 are managed within session domain 720. Both session domains 710 and 720 are managed within session domain 720. Both session domains 710 and 720 are positioned directly beneath the session context 700. In one embodiment, each session domain 710 and 720 contains session objects related to different applications. For example, session domain 710 may contain session objects related to a “shopping cart” application and session domain 720 may contain session objects related to a “calculator” application.

Two additional session domains 730 and 740 are logically positioned beneath session domain 710 within the hierarchy. In one embodiment, these session domains 730 and 740 include session objects 731, 732 and 741, 742, respectively, from the same application associated with session domain 710. Returning to the previous example, if session domain 710 is associated with a shopping cart application, session domains 730 and 740 include session objects related to the shopping cart application.

In one embodiment, a different session context 700 and associated hierarchy of session domains is maintained by session management logic for each layer of the application server. This embodiment will be described using the example illustrated in **Figure 8** in which a plurality of session objects are

managed for two different applications, identified as applications "A" and "B."

The conversational state between client 830 and application A is maintained via three session bean objects 801-803 within the EJB container 842; an HTTP session object 811 within the Web container 840; and two portal session objects 820-821 within the enterprise portal container. Similarly, the conversational state between client 831 and application A is maintained via two session bean objects 804, 806, one HTTP session object 812, and one portal session object 822.

Finally, the conversational state between client 832 and application B is maintained via two session bean objects 807-808, one HTTP session object 813, and one portal session object 823.

In the illustrated example, session bean objects 801 and 804 are instances of a particular session bean, "EJB A," and session bean objects 802 and 806 are instances of another session bean, "EJB B." Session bean objects 803, 807 and 808 are instances of session beans C, D and E, respectively.

In one embodiment, the enterprise portal container 841 is built on top of the Web container 840. As illustrated, a single HTTP session object 811 may be related to multiple portal session objects 821, 821, which store user-specific session data for a plurality of user-specific Web content (e.g., Web pages) generated during the session. When a user logs in to a Web server, for example, multiple user-specific Web pages may be generated in response to client requests and opened in different windows of the client's Web browser. Thus, a single HTTP session object 811 is managed within the Web container

840 and separate portal session data related for each individual window of the user's browser is stored within the portal session objects 820-821. Said another way, the HTTP session object 811 manages session data related to the entire user session with the application server, whereas the portal session objects store session data for specific pages generated during the session (e.g., "My Yahoo" pages on the "Yahoo" Website).

As illustrated in **Figures 9a-c**, each different type of session object shown in **Figure 8** is managed within a session domain under a different context. For example, as illustrated in **Figure 9a**, separate session domains 910 and 920, are maintained for applications A and B, respectively, under session bean context 900. A separate EJB session domain is used to store and manage session instances for each session bean of each application. For example, under the application A session domain 910, session domain 930 associated with EJB A stores session objects 501, 504 which are instances of session bean A; session domain 940 associated with EJB B stores session objects 502, 506 which are instances of session bean B; and session domain 950 associated with EJB C stores session objects 503 which are instances of session bean C. Similarly, under the application B session domain 920, session domain 960 associated with EJB D stores session objects 507 which are instances of session bean D, and session domain 970 associated with EJB E stores session objects 508 which are instances of session bean E. Thus, all session beans are managed under the same session bean context 900 and grouped under application-

specific and bean-specific session domains, thereby simplifying the management and configuration of similar session objects.

As illustrated in **Figure 9b**, separate HTTP session domains 911 and 921 associated with applications A and B, respectively, are managed under an HTTP session context 901. The HTTP session domains store session data associated with each applications' HTTP sessions. For example, HTTP session objects 511 and 512 are stored within the session domain 911 for application A, and HTTP session object 513 is stored within the session domain 921 for application B.

In addition, in one embodiment, illustrated in **Figure 9c**, separate portal session domains 912 and 922 associated with applications A and B, respectively, are managed under an application portal session context 902. The portal session domains store portal session objects containing portal data associated with each application. Specifically, portal session objects 820-822 are stored under the portal session domain 912 for application A, and portal session object 823 is stored under the session domain 922 for application B.

Figure 10 illustrates additional details related to the configuration of each session domain 1000 according to one embodiment of the invention. Specifically, each session domain 1000 includes a set of configuration policy objects 1003, a set of local storage attributes 1002 and a set of global storage attributes 1001.

The configuration policy objects 1003 define the behavior of each session domain. In one embodiment, the configuration policies implemented by the policy objects include, but are not limited to, a “thresholds” policy for setting limits on the number of sessions objects which are stored within each domain. For example, one session object may be joined to one or more session requests based on the thresholds policy. In addition, In one embodiment, a session access policy is implemented which allows the application or other entity which created the session domain (the “domain owner”) to restrict or otherwise influence session access. For example, the domain owner may prevent multithreaded session access via the session access policy. In one embodiment, the configuration policy 1003 further includes a session invalidation policy which defines the behavior of the session object in the case of session invalidation. For example, as described in greater detail below, in one embodiment, transparent session objects are employed under certain conditions. Moreover, inactive sessions may be stored to persistent storage (e.g., the database or file system) and reactivated at a later time in response to subsequent requests.

In one embodiment, a persistence policy is also implemented within the configuration policy objects 1003 to define whether persistent storage should be used and, if so, the particular type of persistent storage that should be used. Session persistence types may include, but are not limited to, in-memory session persistence (i.e., session objects within the domain are stored within the memory of a single process); in-memory replication persistence (i.e., session object state

is stored within a server process memory and is replicated to other server processes); database-based persistence (i.e., the session objects are stored in the database and may be shared across server processes); file system persistence (i.e., sessions are stored within a directory of the file system and can be shared across processes); and cookie-based persistence (i.e., session data is stored within the client in a cookie). It should be noted, however, that the underlying principles of the invention are not limited to any particular set of configuration properties.

Returning to **Figure 10**, the global storage attributes 1001 of each session domain 1000 define attributes of the session domain which are shared across different virtual machines (e.g., different instances of the session domain implemented on different virtual machines). For example, the global session attributes may specify that all sessions within the particular domain 1000 have the same expiration time (e.g., become inactive/invalid after a period of non-responsiveness from a client). By way of another example, EJB application descriptors may be parsed once and shared between each of the EJB domains located under the session bean context 900, and Web application descriptors may be parsed only once and shared between each of the session objects located under the HTTP session context 901. In one embodiment, global storage attributes are used only in a shared memory implementation (i.e., in which virtual machines 321-325 share session objects via shared memory 340-341 as described herein). In one embodiment, global attributes are identified in shared memory using a unique global attribute name.

The local storage attributes 1002 of the session domain define session attributes which are specific to each individual virtual machine (i.e., they are container-specific). For example, if a particular virtual machine relies on a socket connection to a particular remote server, then this information may be provided within the local storage attributes. Local storage attributes may also include specific references to files within the file system/database and specific references to database connections. In an embodiment which does not employ a shared memory implementation, all attributes of the session domain 1000 are stored within the local storage attributes 1002.

The different containers can use the local storage attributes 1002 to keep certain attributes banded to each concrete server process. For example the HttpSession includes the method `javax.servlet.ServletContext` **`getServletContext()`**. However, `ServletContext` can not be shared between different virtual machines because it maintains references to different resources that are local for the concrete server process. Thus, to enable the application to retrieve the `ServletContext` from the HttpSession, the Web container binds the `ServletContext` as a local attribute for the session domain. The `getServletContext()` method is implemented to provide a lookup of this attribute from the `SessionDomain`. Using this technique, the `getServletContext()` will return different objects based on the server process where the session is activated.

Sessions and caches are very important objects which may consume a significant amount of memory. Thus, it is important to have the ability to observe session and cache object networks without the semantics of the represented data. Accordingly, one embodiment of the invention uses the object graph processing techniques described herein for generating session and cache object graphs and extracting associated data. For example, in one embodiment, object graphs are generated to represent the hierarchical session domains described above.

Figure 11 illustrates one embodiment of an architecture for performing these operations. This embodiment includes a session/cache object graph service 1004 for generating object graph data related to the session object network 1102 and a cache object network 1103. In this embodiment, both the session object network 1102 and the cache object network 1103 are executed within a J2EE engine 1110. As in the prior embodiments (described above with respect to **Figure 4**) a local computer system 411 requests object graph information related to the session object network 1102 and/or the cache object network 1103. The session/cache object graph service 1004 extracts information from the session object network 1102 and the cache object network 1103 to generate an object graph representing the networks as well as the other types of data described herein (e.g., memory size, type, each object's identification code, etc). The object graph and other data are then serialized and transmitted over a network to the requesting computer system 411.

Object graph interpreter logic 406 on the requesting computer system 411 interprets the serialized session/cache object graphs and other data to graphically display the results within a the visualization tool 408. As in the prior embodiments, the object graph and other data may be stored within a storage medium 409 such as a file system or database for later retrieval and viewing. The techniques described above are particularly useful when working with session and cache object networks because these networks/objects may tend to become very large and lead to undesirable memory consumption.

D. Using Object Graphs to Analyze Memory Consumption

Every object within an object-oriented computer program consumes a different amount of a computer's memory. Even objects of the same type can consume different amounts. For example, in a rental car reservation system, if two different people book cars, the object for the person with the longest name will consume more memory for its data presentation. Moreover, the fact that objects may reference other objects makes the problem of measuring memory consumption more difficult. In the foregoing example, each "booking" object holds references to the "person" and "car" objects related to the bookings. Consequently, one object (reservation) holds references to two other objects (person and car). This gets even more complicated when two objects are referencing each other. For example, in an implementation for managing human

resources, each “manager” object holds references for its “employer” objects and every “employer” object holds references to its “manager” objects.

In one embodiment of the invention, to address the foregoing situations, the “characteristic” data for each object is identified (which is not referenced directly by other objects) and the memory difference between the “characteristic” data is measured. In one embodiment, the objects and their references are represented as math graph structures, generated as described above.

Figure 12a illustrates an example using the object relationships previously illustrated in **Figure 3**. In this example, an attempt will be made to measure the difference in memory consumed between nodes 301 and 302. Node 302 references node 305 which is characteristic data for it, and node 301 references node 303 that is not referenced directly from node 302. As a result, the memory consumption difference between node 302 and node 301 is the difference between node 305 and node 303 and the difference of the memory taken by node 302 and node 301 themselves.

One embodiment of a method for measuring the memory consumption difference between two nodes is illustrated in **Figure 12c**. At 1201 all incoming references to both nodes are removed. This is indicated in **Figure 12a** by arrows with dotted lines. At 1202, the resulting sub-graphs are built from both nodes, as indicated in **Figure 12b**. At 1203, the sum of the memory consumed by each node in the sub-graph is calculated for both sub-graphs. In one

embodiment, this is calculated by summing the weights of the node and it's children. By way of example, the memory consumption of node 301 is equal to the combined memory consumption of nodes 301, 303 and 304. Similarly, the memory consumption of node 302 is the is equal to the combined memory consumption of nodes 302, 304 and 304. At 1204, the difference between the sums of the memory consumption for nodes 301 and 302 is calculated by subtracting the sums. In other words $\text{delta}(301,302) = \text{nativeSize}(301) + \text{nativeSize}(303) + \text{nativeSize}(304) - \text{nativeSize}(302) - \text{nativeSize}(304) - \text{nativeSize}(305)$, or $\text{delta}(301,302) = \text{nativeSize}(301) + \text{nativeSize}(303) - \text{nativeSize}(302) - \text{nativeSize}(305)$. As mentioned above, in one embodiment, the following weights of the generic types of the nodes are used: boolean = 1 byte; byte = 1 byte; char = 2 bytes; short = 2 bytes; int = 4 bytes; long = 8 bytes; float = 4 bytes; and double = 8 bytes.

Using the foregoing techniques, the difference between the memory consumed by two objects may be determined and, more specifically, the difference between "unique" parts of memory consumed by these objects. In many cases, the object graph is so complex that each node references directly or indirectly (e.g., through a child) any other node. In these cases, while it is not possible to say in general which node takes less memory, the "delta" measure between the two nodes may still be used.

By way of example, in **Figure 3**, node 302 holds memory of node 303 because it has a reference to 302 which has reference to 303. It also holds a

reference to node 305. The delta between these two objects 301 and 302 depends more on the size of object 303 minus the size of object 305. This, if 301 is bigger using the delta function we will examine its graph and determine that it is because of node 303. Consequently, optimizing its memory (for node 303) will lead to optimizing the memory taken by 301 and 302.

A system architecture on which embodiments of the invention may be implemented is illustrated in **Figure 13**. The architecture includes a plurality of application server “instances” 1301 and 1302. The application server instances 1301 and 1302 each include a group of worker nodes 1312-1314 and 1315-1316 (also sometimes referred to herein as “server nodes”), respectively, and a dispatcher 1311 and 1312, respectively. The application server instances 1301, 1302 communicate through a central services instance 1300 using message passing. In one embodiment, the central services instance 1300 includes a locking service and a messaging service (described below). The combination of all of the application server instances 1301 and 1302 and the central services instance 1300 is referred to herein as a “cluster.” Although the following description will focus solely on instance 1301 for the purpose of explanation, the same principles apply to other instances within the cluster.

The worker/server nodes 1312-1314 within instance 1301 provide the business and presentation logic for the network applications supported by the system including, for example, the Web container 211 and the EJB container 201 functionality described herein. Each of the worker nodes 1012-1014 within a

particular instance may be configured with a redundant set of programming logic and associated data, represented as virtual machines 1321-1323 in **Figure 13**.

In one embodiment, the dispatcher 1311 distributes service requests from clients to one or more of the worker nodes 1312-1314 based on the load on each of the servers. For example, in one embodiment, the dispatcher maintains separate queues for each of the worker nodes 1312-1314 in a shared memory 1340. The dispatcher 1311 fills the queues with client requests and the worker nodes 1312-1314 consume the requests from each of their respective queues. The client requests may be from external clients (e.g., browser requests) or from other components/objects within the instance 1301 or cluster.

In one embodiment, the worker nodes 1312-1314 may be Java 2 Enterprise Edition ("J2EE") worker nodes which support Enterprise Java Bean ("EJB") components and EJB containers (at the business layer) and Servlets and Java Server Pages ("JSP") (at the presentation layer). In this embodiment, the virtual machines 1321-1325 implement the J2EE standard (as well as the additional non-standard features described herein). It should be noted, however, that certain high-level features described herein may be implemented in the context of different software platforms including, by way of example, Microsoft .NET platforms and/or the Advanced Business Application Programming ("ABAP") platforms developed by SAP AG, the assignee of the present application.

As indicated in **Figure 13**, a object graph processing service 1330 such as the one described above may be executed across each of the instances 1301, 1302. In one embodiment, the object graph processing service 1330 implements the various object graph processing techniques described herein.

In one embodiment, communication and synchronization between each of the instances 1301, 1302 is enabled via the central services instance 1300. As mentioned above, the central services instance 1300 includes a messaging service and a locking service. The message service allows each of the servers within each of the instances to communicate with one another via a message passing protocol. For example, messages from one server may be broadcast to all other servers within the cluster via the messaging service (e.g., such as the cache configuration messages described below). Alternatively, messages may be addressed directly to specific servers within the cluster (i.e., rather than being broadcast to all servers). In one embodiment, the locking service disables access to (i.e., locks) certain specified portions of configuration data and/or program code stored within a central database 1345. The locking service locks data on behalf of various system components which need to synchronize access to specific types of data and program code. In one embodiment, the central services instance 1300 is the same central services instance as implemented within the Web Application Server version 6.3 and/or 6.4 developed by SAP AG. However, the underlying principles of the invention are not limited to any particular type of central services instance.

In addition, unlike prior systems, one embodiment of the invention shares objects across virtual machines 1321-1325. Specifically, in one embodiment, objects such as session objects which are identified as “shareable” are stored within a shared memory region 1340, 1341 and are made accessible to multiple virtual machines 1321-1325. Creating new object instances from scratch in response to client requests can be a costly process, consuming processing power and network bandwidth. As such, sharing objects between virtual machines as described herein improves the overall response time of the system and reduces server load.

In a shared memory implementation, a shared memory area 1340, 1341 or “heap” is used to store data objects that can be accessed by multiple virtual machines 1321-1325. The data objects in a shared memory heap should generally not have any pointers or references into any private heap (e.g., the private memory regions/heaps of the individual virtual machines). This is because if an object in the shared memory heap had a member variable with a reference to a private object in one particular virtual machine, that reference would be invalid for all the other virtual machines that use that shared object.

More formally, this restriction can be thought of as follows: For every shared object, the transitive closure of the objects referenced by the initial object should only contain shared objects at all times. Accordingly, in one implementation of the invention, objects are not put into the shared memory heap by themselves – rather, objects (such as the scheduling analysis module

402 and scheduling file 404 described above) are put into the shared memory heap in groups known as "shared closures." A shared closure is an initial object plus the transitive closure of all the objects referenced by the initial object.

Embodiments of the invention may include various steps as set forth above. The steps may be embodied in machine-executable instructions which cause a general-purpose or special-purpose processor to perform certain steps. Alternatively, these steps may be performed by specific hardware components that contain hardwired logic for performing the steps, or by any combination of programmed computer components and custom hardware components.

Elements of the present invention may also be provided as a machine-readable medium for storing the machine-executable instructions. The machine-readable medium may include, but is not limited to, flash memory, optical disks, CD-ROMs, DVD ROMs, RAMs, EPROMs, EEPROMs, magnetic or optical cards, propagation media or other type of machine-readable media suitable for storing electronic instructions. For example, the present invention may be downloaded as a computer program which may be transferred from a remote computer (e.g., a server) to a requesting computer (e.g., a client) by way of data signals embodied in a carrier wave or other propagation medium via a communication link (e.g., a modem or network connection).

Throughout the foregoing description, for the purposes of explanation, numerous specific details were set forth in order to provide a thorough understanding of the invention. It will be apparent, however, to one skilled in the

art that the invention may be practiced without some of these specific details.

For example, although many of the embodiments set forth above relate to a Java or J2EE implementation, the underlying principles of the invention may be implemented in virtually any enterprise networking environment. Moreover, although some of the embodiments set forth above are implemented within a shared memory environment, the underlying principles of the invention are equally applicable to a non-shared memory environment.

Accordingly, the scope and spirit of the invention should be judged in terms of the claims which follow.

CLAIMS

What is claimed is:

1. A computer-implemented method comprising:
analyzing relationships between session and/or cache objects within a network of sessions and caches to determine an object network structure for the session and/or cache objects;
generating object graph data representing the object network structure;
serializing the object graph data and transmitting the object graph data over a network to a requesting computer; and
interpreting the object graph data to render a view of the session and/or cache object network structure in a graphical user interface.
2. The computer-implemented method as in claim 1 further comprising:
collecting object data related to each session and/or cache object in the network of objects; and
serializing the object data and transmitting the object data over the network to the requesting computer along with the object network data.
3. The computer implemented method as in any one of the preceding claims further comprising:
displaying the object data related to a session and/or cache object in response to selection of a graphic representing the object on the requesting computer.

4. The computer-implemented method as in any one of the preceding claims wherein the object graph data comprises a list structure.

5. The computer-implemented method as in any one of the preceding claims wherein the object graph data comprises a matrix structure.

6. The computer-implemented method as in any one of the preceding claims in combination with claim 2 wherein the object data comprises a value representing memory consumed by each object.

7. The computer-implemented method as in any one of the preceding claims further comprising:

assigning each object within the network a unique identification code; and
serializing the identification codes and transmitting the identification codes over the network to the requesting computer system.

8. An apparatus having a memory for storing program code and a processor for processing the program code to perform the operations of:

analyzing relationships between session and/or cache objects within a network of sessions and caches to determine an object network structure for the session and/or cache objects;

generating object graph data representing the object network structure;
serializing the object graph data and transmitting the object graph data over a network to a requesting computer; and

interpreting the object graph data to render a view of the session and/or cache object network structure in a graphical user interface.

9. The apparatus as in claim 8 comprising additional program code which, when executed by the processor, performs the operations of:

- collecting object data related to each session and/or cache object in the network of objects; and
- serializing the object data and transmitting the object data over the network to the requesting computer along with the object network data.

10. The computer implemented method as in claim 8 or 9 comprising additional program code which, when executed by the processor, performs the operations of:

- displaying the object data related to a session and/or cache object in response to selection of a graphic representing the object on the requesting computer.

11. The apparatus as in any one of the preceding claims 8 to 10 wherein the object graph data comprises a list structure.

12. The apparatus as in any one of the preceding claims 8 to 11 wherein the object graph data comprises a matrix structure.

13. The apparatus as in any one of the preceding claims in combination with claim 9 wherein the object data comprises a value representing memory consumed by each object.

14. The apparatus as in any one of the preceding claims 8 to 13 comprising additional program code which, when executed by the processor, performs the operations of:

assigning each object within the network a unique identification code; and
serializing the identification codes and transmitting the identification codes
over the network to the requesting computer system.

15. A machine-readable medium having program code stored thereon
which, when executed by a machine, causes the machine to perform the
operations of:

analyzing relationships between session and/or cache objects within a
network of sessions and caches to determine an object network structure for the
session and/or cache objects;

generating object graph data representing the object network structure;
serializing the object graph data and transmitting the object graph data
over a network to a requesting computer; and

interpreting the object graph data to render a view of the session and/or
cache object network structure in a graphical user interface.

16. The machine-readable medium as in claim 15 comprising additional
program code which, when executed by the machine, causes the machine to
perform the operations of:

collecting object data related to each session and/or cache object in the
network of objects; and

serializing the object data and transmitting the object data over the
network to the requesting computer along with the object network data.

17. The machine-readable medium as in claim 15 or 16 comprising
additional program code which, when executed by the processor, performs the
operations of:

displaying the object data related to a session and/or cache object in response to selection of a graphic representing the object on the requesting computer.

18. The machine-readable medium as in any one of the preceding claims 15 to 17 wherein the object graph data comprises a list structure.

19. The machine-readable medium as in any one of the preceding claims 15 to 18 wherein the object graph data comprises a matrix structure.


20. The machine-readable medium as in any one of the preceding claims in combination with claim 16 wherein the object data comprises a value representing memory consumed by each object.


21. The machine-readable medium as in any one of the preceding claims 15 to 20 comprising additional program code which, when executed by the machine, causes the machine to perform the operations of:

assigning each object within the network a unique identification code; and
serializing the identification codes and transmitting the identification codes over the network to the requesting computer system.

22. A computer program product comprising computer-readable instructions which, when loaded and executed on a suitable system perform the steps of a method according to any of the preceding claims 1 to 7.

01 / 16

**Fig. 1a**
(prior art)



03 / 16

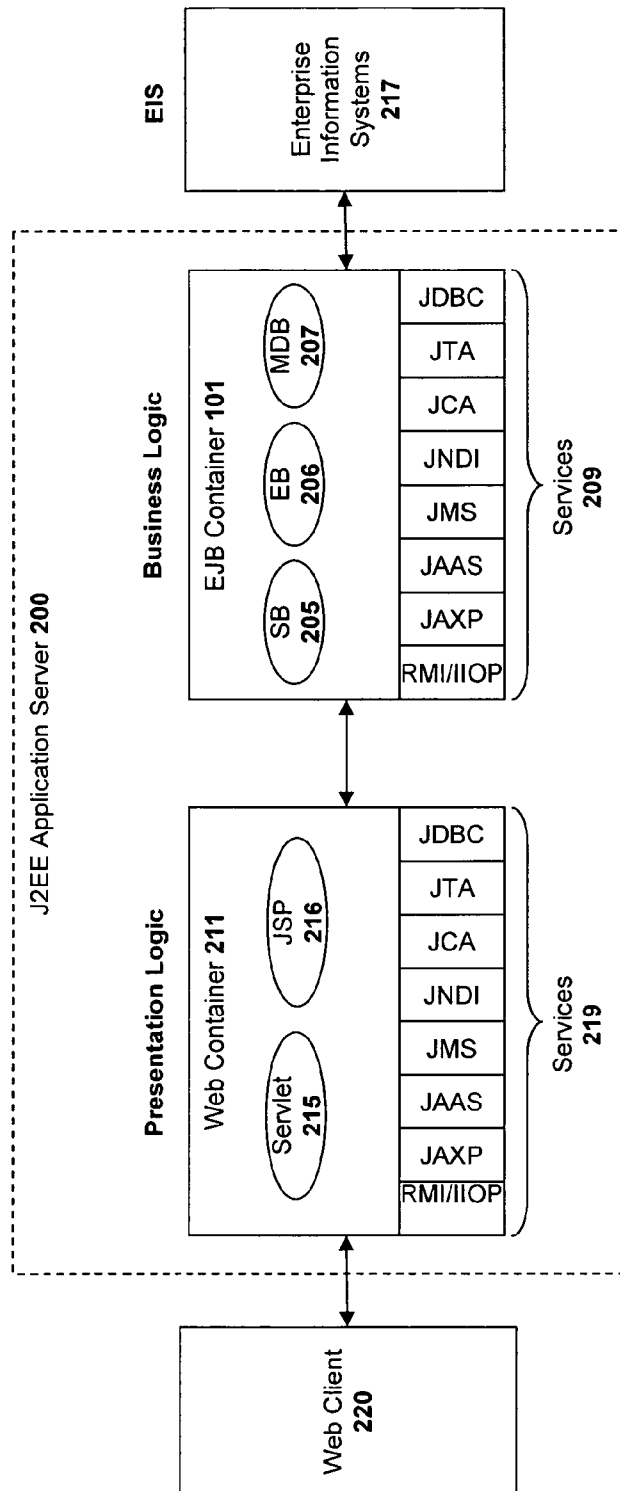




Fig. 2
(prior art)

04 / 16

**Fig. 3**

05 / 16

**Fig. 4**

06/16




Fig. 5

07 / 16




Fig. 6




Fig. 7

09 / 16

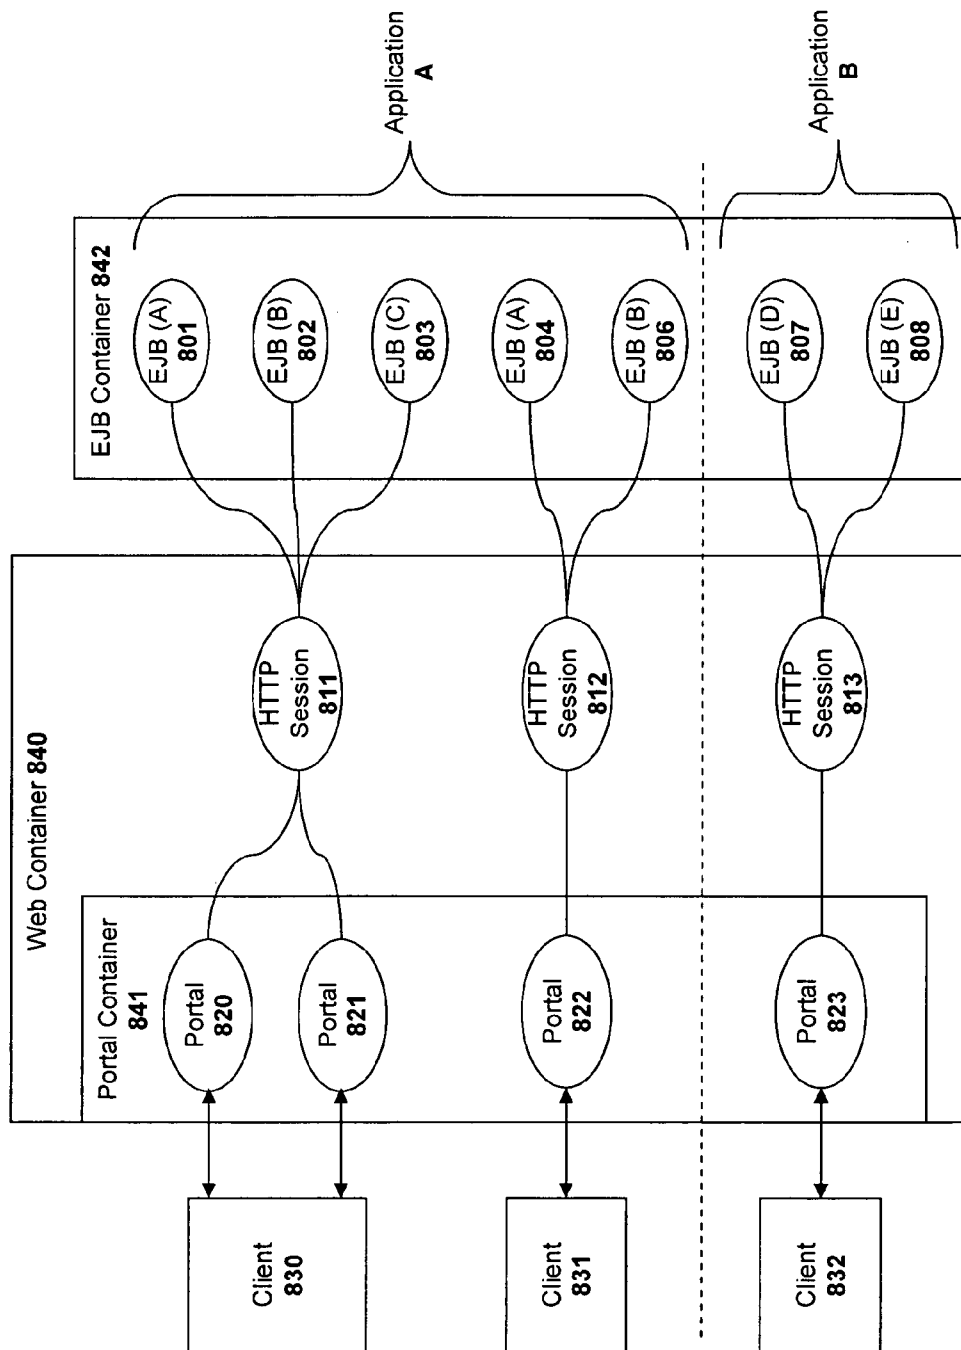


Fig. 8

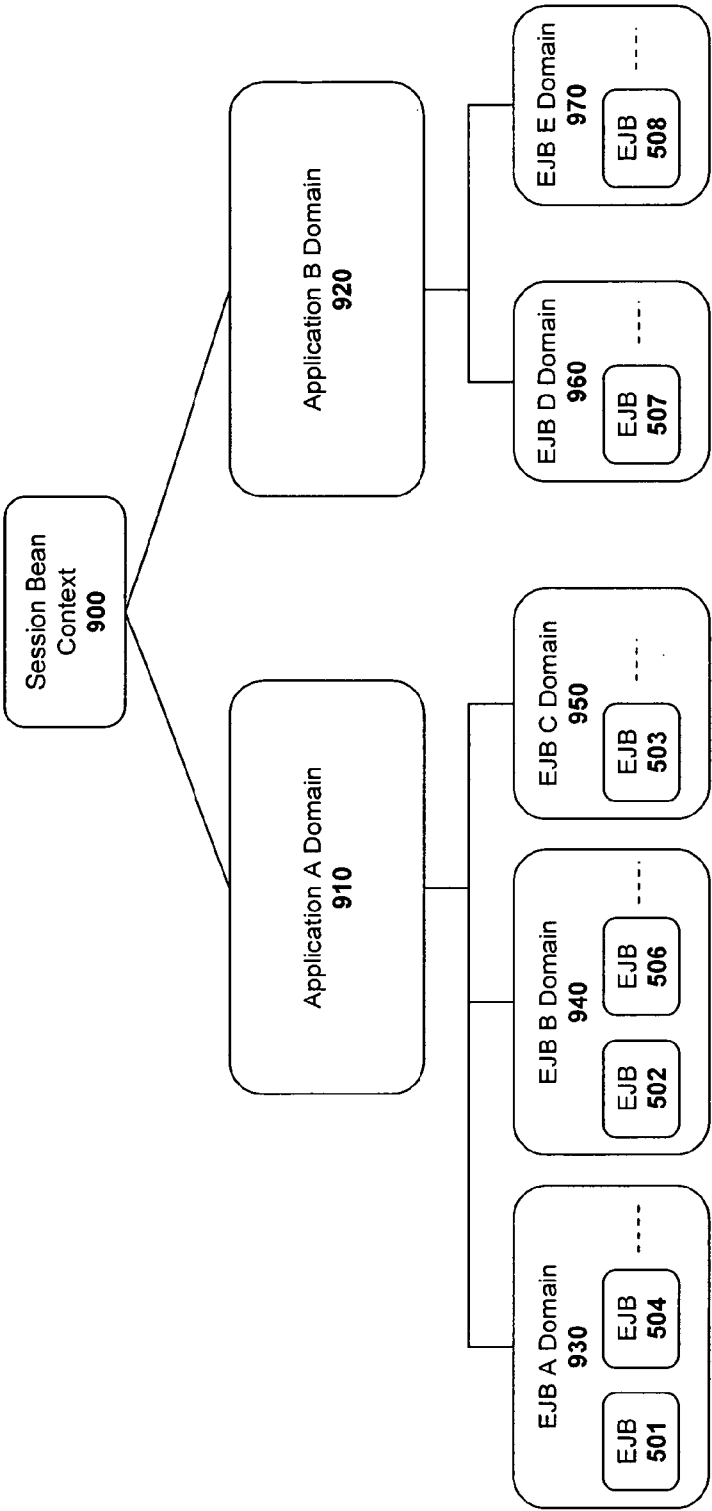



Fig. 9a

11 / 16

**Fig. 9b**




Fig. 10

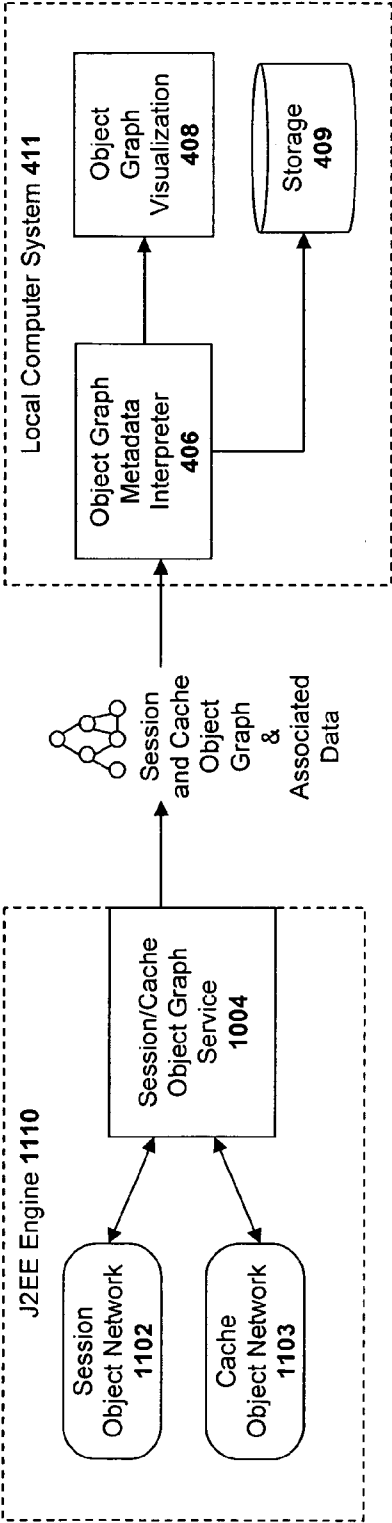


Fig. 11




Fig. 12b


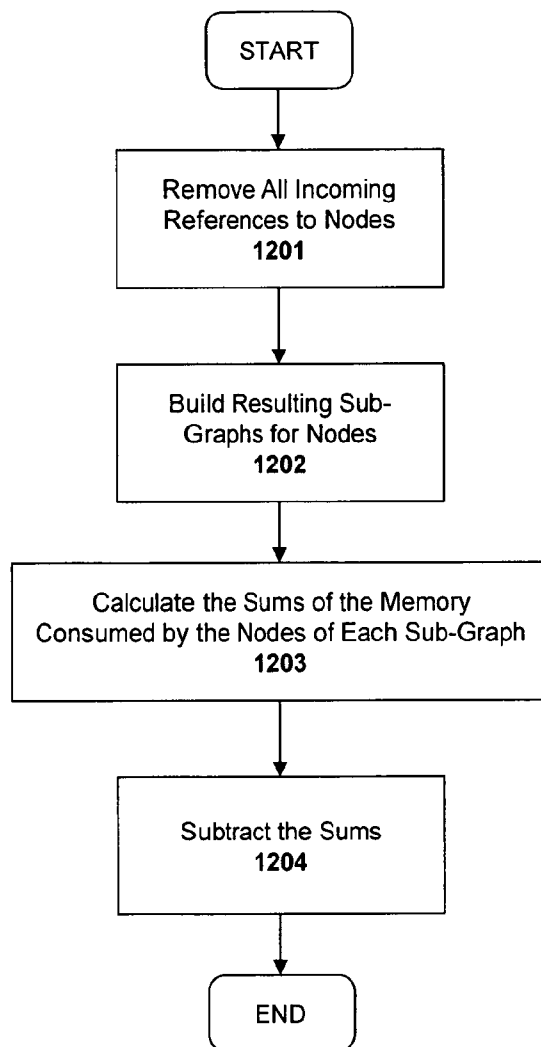


Fig. 12a

15 / 16

**Fig. 12c**

16/16




Fig. 13

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2007/010882

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06F11/30 G06F9/44

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
G06F

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

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>MAHER M. SALAH: "An Environment for Comprehending the Behavior of Software Systems - PhD Thesis" June 2005 (2005-06), DREXEL UNIVERSITY, XP002477233 page 15 - page 40 page 41 - page 60 figure 2.1 page 7 - page 9 figures 3.1,3.2,3.5,3.17</p> <p style="text-align: center;">----- -/-</p>	1-22

☒ Further documents are listed in the continuation of Box C.

☐ See patent family annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- * & * document member of the same patent family

Date of the actual completion of the international search

21 April 2008

Date of mailing of the international search report

07/05/2008

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Schäfer, Andreas

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2007/010882

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	DE PAUW W ET AL: "Visualizing the execution of Java programs" SOFTWARE VISUALIZATION. INTERNATIONAL SEMINAR. REVISED PAPERS (LECTURE NOTES IN COMPUTER SCIENCE VOL.2269) SPRINGER-VERLAG BERLIN, GERMANY, 2002, pages 151-162, XP002477230 ISBN: 3-540-43323-6 Section 1 Section 2 Section 3 figure 2	1-22
A	DE PAUW W ET AL: "Web Services Navigator: visualizing the execution of Web services" IBM SYSTEMS JOURNAL IBM USA, vol. 44, no. 4, 2005, pages 821-845, XP002477231 ISSN: 0018-8670 Section: The Visualization Tool Section Data Collection	1-22
X	SMITH M P ET AL: "Identifying structural features of Java programs by analysing the interaction of classes at runtime" 2005 3RD IEEE INTERNATIONAL WORKSHOP ON VISUALIZING SOFTWARE FOR UNDERSTANDING AND ANALYSIS (IEEE CAT. NO. 05EX1225) IEEE PISCATAWAY, NJ, USA, 2005, pages 108-113, XP002477232 ISBN: 0-7803-9540-9 Section 2 table 1 figures 1,4	1-22
A	SMITH M P ET AL: "Providing a user customisable tool for software visualisation at runtime" FOURTH IASTED INTERNATIONAL CONFERENCE ON VISUALIZATION, IMAGING, AND IMAGE PROCESSING ACTA PRESS ANAHEIM, CA, USA, 2004, pages 135-140, XP002477257 ISBN: 0-88986-454-3 table 1 figures 1-4	1-22

-/--

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2007/010882

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	SMITH M P ET AL: "Runtime visualisation of object oriented software" PROCEEDINGS FIRST INTERNATIONAL WORKSHOP ON VISUALIZING SOFTWARE FOR UNDERSTANDING AND ANALYSIS IEEE COMPUT. SOC LOS ALAMITOS, CA, USA, 2002, pages 81-89, XP002477258 ISBN: 0-7695-1662-9 figure 6 Section 3.2.2	1-22
A	GILBERG R.F., FOROUZAN A.A.: "Data Structures: A Pseudocode Approach with C" 31 May 2006 (2006-05-31), THOMSON COURSE TECHNOLOGY 310340 , XP002477259 page 488 - page 491	4,5,11, 12,18,19