(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
30 November 2006 (30.11.2006)

PCT

(10) International Publication Number
WO 2006/126992 A1

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

(21) International Application Number:
PCT/US2005/018009

(22) International Filing Date: 23 May 2005 (23.05.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
11/133,831 20 May 2005 (20.05.2005) US

(71) Applicant *(for all designated States except US)*: **SUN MICROSYSTEMS, INC.** [US/US]; 4150 Network Circle, Santa Clara, California 95054 (US).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **KAMEN, Yury** [US/US]; 184 Beach Park Blvd., Foster City, California 94404 (US). **ALI, Syed M.** [IN/US]; 1069 Polk Ave., Sunnyvale, California 94086 (US). **ALUR, Deepak** [IN/US]; 45770 Cayuga Court, Fremont, California 94539 (US). **CRUPI, John P.** [US/US]; 4806 Westway Drive, Bethesda, Maryland 20816 (US). **MALKS, Daniel B.** [US/US]; 631 South Stafford Street, Arlington, Virginia 22204 (US).

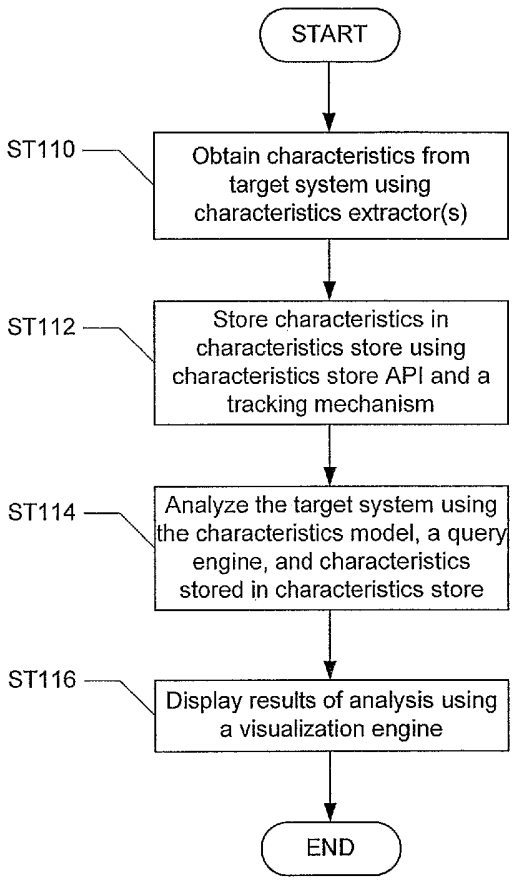
(74) Agent: **OSHA, Jonathan, P.**; Osha Liang L.L.P., 1221 McKinney Street, Suite 2800, Houston, TX 77010 (US).

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

[Continued on next page]

(54) Title: METHOD AND APPARATUS FOR TRACKING CHANGES IN A SYSTEM



(57) Abstract: A method for analyzing a target system (100) that includes obtaining a plurality of characteristics from the target system (100) using a characteristics extractor (102A, 102N), wherein the plurality of characteristics is associated with a characteristics model (108), storing each of the plurality of characteristics in a characteristics store (106) using a tracking mechanism, and analyzing the target system (100) by issuing a query to the characteristics store (106) to obtain an analysis result, wherein the query uses tracking information associated with the tracking mechanism.

WO 2006/126992 A1

European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, MC, 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).

Declarations under Rule 4.17:

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

— *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

— *with international search report*

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.

METHOD AND APPARATUS FOR TRACKING CHANGES IN A SYSTEM

BACKGROUND

[0001] As software technology has evolved, new programming languages and increased programming language functionality has been provided. The resulting software developed using this evolving software technology has become more complex. The ability to manage the quality of software applications (including design quality and architecture quality) is becoming increasingly more difficult as a direct result of the increasingly complex software. In an effort to manage the quality of software applications, several software development tools and approaches are now available to aid software developers in managing software application quality. The following is a summary of some of the types of quality management tools currently available.

[0002] One common type of quality management tool is used to analyze the source code of the software application to identify errors (or potential errors) in the source code. This type of quality management tool typically includes functionality to parse the source code written in a specific programming language (*e.g.,* Java™, C++, etc.) to determine whether the source code satisfies one or more coding rules (*i.e.,* rules that define how source code in the particular language should be written). Some quality management tools of the aforementioned type have been augmented to also identify various coding constructs that may result in security or reliability issues. While the aforementioned type of quality management tools corrects coding errors, it does not provide the software developer with any functionality to verify the quality of the architecture of software application.

[0003] Other quality management tools of the aforementioned type have been augmented to verify that software patterns have been properly implemented. Specifically, some quality management tools of the aforementioned type have been augmented to allow the software developer to indicate, in the source code, the type of software pattern the developer is using. Then the quality management tool verifies, during compile time, that the software pattern was used/implemented correctly.

[0004] In another implementation of the aforementioned type of quality management tools, the source code of the software is parsed and the components (*e.g.,* classes, interfaces, etc.) extracted from the parsing are subsequently combined in a relational graph (*i.e.,* a graph linking all (or sub-sets) of the components). In a subsequent step, the software developer generates an architectural design, and then compares the architectural design to the relational graph to determine whether the software application conforms to the architectural pattern. While the aforementioned type of quality management tool enables the software developer to view the relationships present in the software application, it does not provide the software developer with any functionality to conduct independent analysis on the extracted components.

[0005] Another common type of quality management tool includes functionality to extract facts (*i.e.,* relationships between components (classes, interfaces, etc.) in the software) and subsequently displays the extracted facts to the software developer. While the aforementioned type of quality management tool enables the software developer to view the relationships present in the software application, it does not provide the developer with any functionality to independently query the facts or any functionality to extract information other than facts from the software application.

[0006] Another common type of quality management tool includes functionality to extract and display various statistics (*e.g.,* number of lines of code, new artifacts added, software packages present, etc.) of the software application to the software developer. While the aforementioned type of quality management tool enables the software developer to view the current state of the software application, it does not provide the developer with any functionality to verify the quality of the architecture of the software application.

SUMMARY

[0007] In general, in one aspect, the invention relates to a method for analyzing a target system, comprising obtaining a plurality of characteristics from the target system using a characteristics extractor, wherein the plurality of characteristics is associated with a characteristics model, storing each of the plurality of characteristics in a characteristics store using a tracking mechanism, and analyzing the target system by issuing a query to the characteristics store to obtain an analysis result, wherein the query uses tracking information associated with the tracking mechanism.

[0008] In general, in one aspect, the invention relates to a system, comprising a characteristics model defining at least one artifact and a plurality of characteristics associated with the at least one artifact, a target system comprising at least one of the plurality of characteristics defined in the characteristics model, at least one characteristics extractor configured to obtain at least one of the plurality of characteristics from the target system, a characteristics store configured to store the at least one of the plurality of characteristics obtained from the target system using a tracking mechanism, and a query engine configured to analyze the target system by issuing a query to the characteristics store and configured to obtain an analysis result in response to the at least one

query, wherein the query uses tracking information associated with the tracking mechanism.

[0009] In general, in one aspect, the invention relates to a computer readable medium comprising software instructions for analyzing a target system, comprising software instructions to obtain a plurality of characteristics from the target system using a characteristics extractor, wherein the plurality of characteristics is associated with a characteristics model, store each of the plurality of characteristics in a characteristics store using a tracking mechanism, and analyze the target system by issuing a query to the characteristics store to obtain an analysis result, wherein the query uses tracking information associated with the tracking mechanism.

[0010] Other aspects of the invention will be apparent from the following description and the appended claims.

BRIEF DESCRIPTION OF DRAWINGS

[0011] FIG. 1 shows a system in accordance with one embodiment of the invention.

[0012] FIG. 2 shows a characteristic storage format in accordance with one embodiment of the invention.

[0013] FIG. 3 shows a characteristics store in accordance with one embodiment of the invention.

[0014] FIG. 4 shows a characteristics model in accordance one embodiment of the invention.

[0015] FIGS. 5 and 6 show flowcharts in accordance with one embodiment of the invention.

[0016] FIG. 7 shows a computer system in accordance with one embodiment of the invention.

DETAILED DESCRIPTION

[0017] Exemplary embodiments of the invention will be described with reference to the accompanying drawings. Like items in the drawings are shown with the same reference numbers.

[0018] In the exemplary embodiment of the invention, numerous specific details are set forth in order to provide a more thorough understanding of the invention. However, it will be apparent to one of ordinary skill in the art that the invention may be practiced without these specific details. In other instances, well-known features have not been described in detail to avoid obscuring the invention.

[0019] In general, embodiments of the invention relate to a method and apparatus for pattern-based system design analysis. More specifically, embodiments of the invention provide a method and apparatus for using one or more characteristics models, one or more characteristics extractors, and a query engine configured to query the characteristics of a target system to analyze the system design. Embodiments of the invention provide the software developer with a fully configurable architectural quality management tool that enables the software developer to extract information about the characteristics of the various artifacts in the target system, and then issue queries to determine specific details about the various artifacts including, but not limited to, information such as: number of artifacts of the specific type present in the target system, relationships between the various artifacts in the target system, the interaction of the various artifacts within the target system, the patterns that are used within the target system, etc.

[0020] Further, embodiments of the invention provide a method and apparatus for tracking changes in a target system. More specifically, embodiments of the

invention provide tracking mechanisms to track changes in the target system. The changes in the target system may include, but are not limited to, adding/removing/modifying artifacts from the target system, adding/removing/modifying characteristics within the target system, and adding/removing/modifying relationships within the system, etc.

[0021] FIG. 1 shows a system in accordance with one embodiment of the invention. The system includes a target system (100) (*i.e.,* the system that is to be analyzed) and a number of components used in the analysis of the target system. In one embodiment of the invention, the target system (100) may correspond to a system that includes software, hardware, or a combination of software and hardware. More specifically, embodiments of the invention enable a user to analyze specific portions of a system or the entire system. Further, embodiments of the invention enable a user to analyze the target system with respect to a specific domain (discussed below). Accordingly, the target system (100) may correspond to any system under analysis, where the system may correspond to the entire system including software and hardware, or only a portion of the system (*e.g.,* only the hardware portion, only the software portion, a sub-set of the hardware or software portion, or any combination thereof). As shown in FIG. 1, the system includes the following components to aid in the analysis of the target system: one or more characteristics extractors (*e.g.,* characteristics extractor A (102A), characteristics extractor N (102N)), a characteristics store application programming interface (API) (104), a characteristics store (106), a characteristics model (108), a query engine (110), and visualization engine (112). Each of these components is described below.

[0022] In one embodiment of the system, the characteristics model (108) describes artifacts (*i.e.,* discrete components) in a particular domain. In one embodiment of the invention, the domain corresponds to any grouping of "related artifacts" (*i.e.,* there is a relationship between the artifacts). Examples of domains include,

but are not limited to, a Java™ 2 Enterprise Edition (J2EE) domain (which includes artifacts such as servlets, filters, welcome file, error page, etc.), a networking domain (which includes artifacts such as web server, domain name server, network interface cards, etc), and a DTrace domain (described below). In one embodiment of the invention, each characteristics model includes one or more artifacts, one or more relationships describing the interaction between the various artifacts, and one or more characteristics that describe various features of the artifact. An example of a characteristics model (108) is shown in Figure 4. Those skilled in the art will appreciate that the system may include more than one characteristics model (108).

[0023] In one embodiment of the invention, the use of a characteristics model (108) enables a user to analyze the target system (100) with respect to a specific domain. Further, the use of multiple characteristics models allows the user to analyze the target system (100) across multiple domains. In addition, the use of multiple characteristics models allows the user to analyze the interaction between various domains on the target system (100).

[0024] In one embodiment of the invention, the characteristics extractors (*e.g.,* characteristics extractor A (102A), characteristics extractor N (102N)) are used to obtain information about various artifacts (*i.e.,* characteristics) defined in the characteristics model (108). In one embodiment of the invention, the characteristics extractors (characteristics extractor A (102A), characteristics extractor B (102N)) are generated manually using the characteristics model (108).

[0025] In one embodiment of the invention, the characteristics extractor (*e.g.,* characteristics extractor A (102A), characteristics extractor B (102N)) corresponds to an agent loaded on the target system (100) that is configured to monitor and obtain information about the artifacts in the target system (100).

Alternatively, the characteristics extractor (*e.g.*, characteristics extractor A (102A), characteristics extractor B (102N)) may correspond to an interface that allows a user to manually input information about one or more artifacts in the target system (100). In another embodiment of the invention, the characteristics extractor (*e.g.*, characteristics extractor A (102A), characteristics extractor B (102N)) may correspond to a process (or system) configured to obtain information about one or more artifacts in the target system (100) by monitoring network traffic received by and sent from the target system (100). In another embodiment of the invention, the characteristics extractor (*e.g.*, characteristics extractor A (102A), characteristics extractor B (102N)) may correspond to a process (or system) configured to obtain information about one or more artifacts in the target system (100) by sending requests (*e.g.*, pinging, etc.) for specific pieces of information about artifacts in the target system (100) to the target system (100), or alternatively, sending requests to the target system and then extracting information about the artifacts from the responses received from target system (100). Those skilled in the art will appreciate that different types of characteristics extractors may be used to obtain information about artifacts in the target system (100).

[0026] Those skilled in the art will appreciate that each characteristics extractor (or set of characteristics extractors) is associated with a particular characteristics model (108). Thus, each characteristics extractor typically only retrieves information about artifacts described in the characteristics model with which the characteristics extractor is associated. Furthermore, if there are multiple characteristics models in the system, then each characteristics model may be associated with one or more characteristics extractors.

[0027] The information about the various artifacts in the target system (100) obtained by the aforementioned characteristics extractors (*e.g.*, characteristics extractor A (102A), characteristics extractor N (102N)) is stored in the

characteristics store (106) via the characteristic store API (104). In one embodiment of the invention, characteristics store API (104) provides an interface between the various characteristics extractors (characteristics extractor A (102A), characteristics extractor N (102N)) and the characteristics store (106). Further, the characteristics store API (104) includes information about where in the characteristics store (106) each characteristic obtained from the target system (100) should be stored.

[0028] In one embodiment of the invention, the characteristics store (106) corresponds to any storage that includes functionality to store characteristics in a manner that allows the characteristics to be queried. In one embodiment of the invention, the characteristics store (106) may correspond to a persistent storage device (*e.g.,* hard disk, etc). In one embodiment of the invention, the characteristics store (106) corresponds to a relational database that may be queried using a query language such as Structure Query Language (SQL). Those skilled in the art will appreciate that any query language may be used. In one embodiment of the invention, if the characteristics store (106) is a relational database, then the characteristics store (106) includes a schema associated with the characteristics model (108) that is used to store the characteristics associated with the particular characteristics model (108). Those skilled in the art will appreciate that, if there are multiple characteristics models, then each characteristics model (108) may be associated with a separate schema.

[0029] In one embodiment of the invention, if the characteristics store (106) is a relational database that includes a schema associated with the characteristics model (108), then the characteristics store API (104) includes the necessary information to place characteristics obtained from target system (100) in the appropriate location in the characteristics store (106) using the schema.

[0030] In one embodiment of the invention, the query engine (110) is configured to issue queries to the characteristics store (106). In one embodiment of the invention, the queries issued by the query engine (110) enable a user (*e.g.,* a system developer, etc.) to analyze the target system (100). In particular, in one embodiment of the invention, the query engine (110) is configured to enable the user to analyze the presence of specific patterns in the target system as well as the interaction between various patterns in the target system.

[0031] In one embodiment of the invention, the query engine (110) includes functionality to issue queries using a tracking mechanism. More specifically, the query engine (110) includes functionality to issue a query to the characteristics store (106) that includes tracking information associated with the tracking mechanism. More specifically, the query engine (110) allows a user to determine changes in the system overtime by issuing queries that leverage tracking information associated with the tracking mechanism(s) used to store the characteristics in the characteristics store (106). Examples of such queries are discussed below.

[0032] In one embodiment of the invention, a pattern corresponds to a framework that defines how specific components in the target system (100) should be configured (*e.g.,* what types of information each component should manage, what interfaces should each component expose), and how the specific components should communicate with each other (*e.g.,* what data should be communicated to other components, etc.). Patterns are typically used to address a specific problem in a specific context (*i.e.,* the software/system environment in which the problem arises). Said another way, patterns may correspond to a software architectural solution that incorporates best practices to solve a specific problem in a specific context.

[0033] Continuing with the discussion of FIG. 1, the query engine (110) may also be configured to issue queries about interaction of specific patterns with components that do not belong to a specific pattern. Further, the query engine (110) may be configured to issue queries about the interaction of components that do not belong to any patterns.

[0034] In one embodiment of the invention, the query engine (110) may include pre-specified queries and/or enable to the user to specify custom queries. In one embodiment of the invention, both the pre-specified queries and the custom queries are used to identify the presence of one or more patterns and/or the presence of components that do not belong to a pattern in the target system (100). In one embodiment of the invention, the pre-specified queries and the custom queries are specified using a Pattern Query Language (PQL). In one embodiment of the invention, PQL enables the user to query the artifacts and characteristics of the artifacts stored in the characteristics store (106) to determine the presence of a specific pattern, specific components of a specific pattern, and/or other components that are not part of a pattern, within the target system (100).

[0035] In one embodiment of the invention, the query engine (110) may include information (or have access to information) about the characteristics model (108) that includes the artifact and/or characteristics being queried. Said another way, if the query engine (110) is issuing a query about a specific artifact, then the query engine (110) includes information (or has access to information) about the characteristics model to which the artifact belongs. Those skilled in the art will appreciate that the query engine (110) only requires information about the particular characteristics model (108) to the extent the information is required to issue the query to the characteristics store (106).

[0036] Those skilled in the art will appreciate that the query engine (110) may include functionality to translate PQL queries (*i.e.,* queries written in PQL) into queries written in a query language understood by the characteristics store (106) (*e.g.,* SQL). Thus, a query written in PQL may be translated into an SQL query prior to being issued to the characteristics store (106). In this manner, the user only needs to understand the artifacts and/or characteristics that the user wishes to search for and how to express the particular search using PQL. The user does not need to be concerned with how the PQL query is handled by the characteristics store (106).

[0037] Further, in one or more embodiments of the invention, PQL queries may be embedded in a programming language such as Java™, Groovy, or any other programming language capable of embedding PQL queries. Thus, a user may embed one or more PQL queries into a program written in one of the aforementioned programming languages. Upon execution, the program issues one or more PQL queries embedded within the program and subsequently receives and processes the results prior to displaying them to the user. Those skilled in the art will appreciate that the processing of the results is performed using functionality of the programming language in which the PQL queries are embedded.

[0038] In one embodiment of the invention, the results of the individual PQL queries may be displayed using the visualization engine (112). In one embodiment of the invention, the visualization engine (112) is configured to output the results of the queries on a display device (*i.e.*, monitor, printer, projector, etc.).

[0039] In one embodiment of the invention, the system may include tracking mechanisms to tracking changes in the target system. More specifically, the system includes functionality to use a tracking mechanism to store characteristics

within the system. In one embodiment of the invention, the tracking mechanism corresponds to associating each characteristic stored in the characteristics store (106) with a time-related attribute. (*See* FIG. 2) Alternatively, the tracking mechanism corresponds to using sub-characteristics stores to store characteristics associated with an instance of the target system (*e.g.,* a particular version, the target system at a given time, etc.). (*See* FIG. 3).

[0040] In one embodiment of the invention, the query engine (110) is configured to issue queries to the characteristics store (106) using tracking information associated with the tracking mechanisms used to the store the characteristics within the characteristics store (106). For example, if the characteristics are stored with a time-related attribute, then the query engine may issue queries that specify a particular timeframe, version, etc. Thus, only characteristics that have a time-related attributed within the range specified by the query are considered. Alternatively, if sub-characteristics stores are used, then the query may specify a particular sub-characteristics store (*e.g.,* queries that specify the name of the particular sub-characteristics store) within the characteristics store (106) to issue the queries to.

[0041] FIG. 2 shows a characteristic storage format in accordance with one embodiment of the invention. As shown in FIG. 2, the characteristic (114) is associated with a time-related attribute (116). In one embodiment of the invention, the time-related attribute (116) may include, but is not limited to, a version of the target system (or a project) and a timestamp. In one embodiment of the invention, if the time-related attribute (116) corresponds to a the version of the target system (or the project), then the version may correspond, but is not limited to, the version of the target system (or project) at the time the particular characteristic was created, the version of the target system (or project) at the time the characteristic was last modified, or the version of the target system (or project) at the time the characteristic was extracted from the target system and

stored in the characteristics store. In one embodiment of the invention, if the time-related attribute (116) corresponds to the timestamp, then the timestamp may correspond, but is not limited to, the timestamp at the time the particular characteristic was created, the timestamp at the time the characteristic was last modified, or the timestamp at the time the characteristic was extracted from the target system and stored in the characteristics store. Those skilled in the art will appreciate that the time-related attribute for a given characteristic does not necessarily need to be stored with the characteristic (*i.e.,* in the same table as the characteristic). Rather, the time-related attribute for a given characteristics must only be stored in the characteristics store (106) or in a location accessible by the system.

[0042] FIG. 3 shows a characteristics store in accordance with one embodiment of the invention. As shown in FIG. 3, the characteristics store (106) includes a number of sub-characteristics stores (118A, 118N). In one embodiment of the invention, each sub-characteristics store (118A, 118N) includes characteristics of the target system obtained during a given timeframe, at a given time, or for a particular version of the target system. Because the sub-characteristics store (118A, 118N) are associated with a version, time, or a timeframe, there is typically no requirement to store a time-related attribute with each of the characteristics in a given sub-characteristics store (118A, 118N). In one embodiment of the invention, the timeframe corresponds to a discrete time (*e.g.,* second, minute, hour, day, week, etc.) during which the characteristics in the particular sub-characteristics store (118A, 118N) were extracted from the target system. In one embodiment of the invention, the sub-characteristics store (118A, 118N) stores a snapshot (*i.e.,* characteristics of the target system at a specific time). Those skilled in the art will appreciate that both tracking mechanisms may be used within a given system.

[0043] As discussed above, each characteristics model defines one or more artifacts, one or more relationships between the artifacts, and one or more characteristics for each artifact. The following is an example of a DTrace characteristics model. In the example, the DTrace characteristics model includes the following attributes: DTraceProject, Network, Computers, CPUs, Processes, Threads, Callstacks, and FunctionCalls. The DTrace characteristics model defines the following relationships between the aforementioned artifacts: DTraceProject includes one or more Networks, each Network includes one or more Computer, each Computer includes one or more CPUs, each CPU runs (includes) one or more Processes, each Process includes one or more Threads, each Thread includes one or more CallStacks, and each CallStacks includes one or more FunctionCalls.

[0044] The following characteristics are used in the DTrace characteristics model: id (*i.e.,* unique CPU id), probeTimestamp (*i.e.,* the performance probe timestamp), memoryCapacity (*i.e.,* the memory available to artifact), cpuNumber (*i.e.,* the number of this CPU in the Computer), usagePercentIO (*i.e.,* the total IO usage percent), usagePercentCPU (*i.e.,* the total CPUusage percent), usagePercentMemory (*i.e.,* the total memory usage percent), usagePercentNetwork (*i.e.,* the total network bandwidth usage percent), usagePercentIOKernel (*i.e.,* the kernel IO usage percent), UsagePercentCPUKernel (*i.e.,* the kernel CPUusage percent), UsagePercentMemoryKernel (*i.e.,* the kernel memory usage percent), and usagePercentNetworkKernel (*i.e.,* the kernel network bandwidth usage percent).

[0045] The following is a DTrace characteristics model in accordance with one embodiment of the invention.

[0046]

DTrace Characteristics Model

```
persistent class DTraceProject {
Long id;
Timestamp probeTimestamp;
String name;
owns Network theNetworks(0,n) inverse theDTraceProject(1,1);
} // class DTraceProject

persistent class Computer {
Long id;
Timestamp probeTimestamp;
String name;
Long numberOfCPUs;
Long memoryCapacity;
Float usagePercentIO;
Float usagePercentCPU;
Float usagePercentMemory;
Float usagePercentNetwork;
Float usagePercentIOKernel;
Float usagePercentCPUKernel;
Float usagePercentMemoryKernel;
Float usagePercentNetworkKernel;
owns CPU theCPUs(0,n) inverse theComputer(1,1);
} // class Computer

persistent class CPU {
Long id;
Timestamp probeTimestamp;
Long cpuNumber;
Long memoryCapacity;
Float usagePercentIO;
Float usagePercentCPU;
Float usagePercentMemory;
Float usagePercentNetwork;
Float usagePercentIOKernel;
Float usagePercentCPUKernel;
Float usagePercentMemoryKernel;
Float usagePercentNetworkKernel;
owns Process theProcesss(0,n) inverse theCPU(1,1);
} // class CPU

persistent class Network {
Long id;
```

```
Timestamp probeTimestamp;
String name;
Long totalCapacity;
Float usagePercent;
owns Computer theComputers(0,n) inverse theNetwork(1,1);
} // class Network

persistent class Process {
Long id;
Timestamp probeTimestamp;
String name;
String commandLine;
Integer priority;
owns Thread theThreads(0,n) inverse theProcess(1,1);
references Process theProcesss(0,n) inverse theProcess(1,1);
} // class Process

persistent class CallStack {
Long id;
Timestamp probeTimestamp;
Float usagePercentIO;
Float usagePercentCPU;
Float usagePercentMemory;
Float usagePercentNetwork;
Float usagePercentIOKernel;
Float usagePercentCPUKernel;
Float usagePercentMemoryKernel;
Float usagePercentNetworkKernel;
owns FunctionCall theFunctionCalls(0,n) inverse theCallStack(1,1);
} // class CallStack

persistent class Thread {
Long id;
String name;
Timestamp probeTimestamp;
Long priority;
Float usagePercentIO;
Float usagePercentCPU;
Float usagePercentMemory;
Float usagePercentNetwork;
Float usagePercentIOKernel;
Float usagePercentCPUKernel;
Float usagePercentMemoryKernel;
Float usagePercentNetworkKernel;
owns CallStack theCallStacks(0,n) inverse theThread(1,1);
```

```
} // class Thread

persistent class FunctionCall {
Long id;
String name;
Timestamp probeTimestamp;
Float usagePercentIO;
Float usagePercentCPU;
Float usagePercentMemory;
Float usagePercentNetwork;
Float usagePercentIOKernel;
Float usagePercentCPUKernel;
Float usagePercentMemoryKernel;
Float usagePercentNetworkKernel;
references FunctionCall theFunctionCalls(0,n) inverse theFunctionCall(1,1);
} // class FunctionCall
```

[0047] In the above DTrace Characteristics Model, the DTraceProject artifact is defined in lines 1-6, the Network artifact defined in lines 41-48, the Computer artifact is defined in lines 8-23, the CPU artifact is defined in lines 25-39, the Processes artifact is defined in lines 50-58, the Thread artifact is defined in lines 74-88, the Callstacks artifact is defined in 61-72, and the FunctionCall artifacts is defined in lines 90-103.

[0048] A graphical representation of the aforementioned DTrace characteristics model is shown in FIG 4. Specifically, the graphical representation of the DTrace characteristics model shows each of the aforementioned artifacts, characteristics associated with each of the aforementioned artifacts, and the relationships (including cardinality) among the artifacts. In particular, box (120) corresponds to the DTraceProject artifact, box (122) corresponds to the Network artifact, box (124) corresponds to the Computer artifact, box (126) corresponds to the CPU artifact, box (128) corresponds to the Process artifact, box (130) corresponds to the Thread artifact, box (132) corresponds to the CallBack artifact, and box (134) corresponds to the FunctionCall artifact.

[0049] FIG. 5 shows a flowchart in accordance with one embodiment of the invention. Initially, a characteristics model is obtained (ST100). In one embodiment of the invention, the characteristics model is obtained from a pre-defined set of characteristics models. Alternatively, the characteristics model is customized characteristics model to analyze a specific domain in the target system and obtained from a source specified by the user.

[0050] Continuing with the discussion of FIG. 5, a schema for the characteristics store is subsequently created and associated with characteristics model (ST102). One or more characteristics extractors associated with characteristics model are subsequently created (ST104). Finally, a characteristics store API is created (ST106). In one embodiment of the invention, creating the characteristics store API includes creating a mapping between characteristics obtained by the characteristics extractors and tables defined by the schema configured to store the characteristics in the characteristics store.

[0051] Those skilled in the art will appreciate that ST100-ST106 may be repeated for each characteristics model. In addition, those skilled in the art will appreciate that once a characteristics store API is created, the characteristics store API may only need to be modified to support additional schemas in the characteristics data store and additional characteristics extractors. Alternatively, each characteristics model may be associated with a different characteristics store API. Those skilled in the art will appreciate that a determination may be made during the setup of the system as to the tracking mechanism to use. Once the tracking mechanism is selected, the appropriate modifications are made to the system. These changes may include modifications to the characteristics store API(s) to store the extracted characteristics in the appropriate sub-characteristics store, modifications to the characteristics extractors to extract additional information from the target system (*e.g.,* time-related attribute information or information used to determine the particular sub-characteristics store a particular

characteristic is to be stored), changes to the schema (as well as the underlying implementation in the characteristics store) to support the storage of time-related attributes and/or sub-characteristics stores.

[0052] At this stage, the system is ready to analyze a target system. FIG. 6 shows a flowchart in accordance with one embodiment of the invention. Initially, characteristics are obtained from the target system using one or more characteristics extractors (ST110). In one embodiment of the invention, the characteristics extractors associated with a given characteristics model only obtain information about characteristics associated with the artifacts defined in the characteristics model.

[0053] Continuing with the discussion of FIG. 6, the characteristics obtained from the target system using the characteristics extractors are stored in the characteristics store using the characteristics store API and a tracking mechanism (ST112). Depending on the tracking mechanism, the characteristics extracted from the target system may be stored in the appropriate sub-characteristics store and/or with a time-related attribute. Those skilled in the art will appreciate that the characteristics extractors will include functionality to obtain the necessary information (*e.g.,* time, timestamp, version, etc.) as required by the tracking mechanism.

[0054] Continuing with the discussion of FIG. 6, once the characteristics are stored in the characteristics store, the target system may be analyzed using the characteristics model (or models), a query engine, and the characteristics stored in the characteristics store (ST114). In one embodiment of the invention, the user uses the query engine to issue queries to characteristics store. As discussed above, the query engine may include information (or have access to information) about the characteristics models currently being used to analyze the target system. In addition, the query may use the tracking information associated with

the tracking mechanisms to track changes in the target system. The results of the analysis are subsequently displayed using a visualization engine (ST116).

[0055] Those skilled in the art will appreciate that ST110-ST112 may be performed concurrently with ST114-ST116. In addition, steps in FIG. 5, may be performed concurrently with the steps in FIG. 6.

[0056] As discussed above, the query engine may issue queries that use the tracking information to analyze changes within the target system. More specifically, the query engine may issue queries that use the time-related attributes or the sub-characteristics stores to track the changes within the target system. The following are two examples of queries that use the time-related attribute to track changes within the target system:

Example 1: Query using Time-Related Attribute

```
SELECT c.name FROM JClass c WHERE c.system.time_stamp > '2005-04-01 10:00'
```

[0057] The above query is used to determine all Java™ classes added to the system since 2005-04-01 10:00.

Example 2: Query using Time-Related Attribute

```
SELECT c.name FROM JClass c WHERE c.system.version > '5.1'
```

[0058] The above query is used to determine all Java™ classes added to the system since version 5.1.

[0059] The following is an example of a query that uses the sub-characteristics stores to track changes within the target system:

Example 3: Query using Sub-Characteristics Stores

```
SELECT c.name FROM classes branch1.c
WHERE (c instanceof "javax.servlet.http.HttpServlet"
    OR c instanceof "javax.servlet.http.GenericServlet")
  AND branch1.c NOT IN
```

```
(SELECT c.name FROM classes branch2.c
 WHERE  c instanceof "javax.servlet.http.HttpServlet"
   OR c instanceof "javax.servlet.http.GenericServlet")
```

[0060] The above query is used to determine all servlets that are in sub-characteristics store 2 (denoted as branch1 in the query) and not in sub-characteristics store 1 (denoted as branch2 in the query). Thus, the above query may be used to determine which servlets were added to the target system between the time sub-characteristics store 1 was populated and the time the sub-characteristics store 2 was populated.

[0061] An embodiment of the invention may be implemented on virtually any type of computer regardless of the platform being used. For example, as shown in FIG. 7, a networked computer system (200) includes a processor (202), associated memory (204), a storage device (206), and numerous other elements and functionalities typical of today's computers (not shown). The networked computer (200) may also include input means, such as a keyboard (208) and a mouse (210), and output means, such as a monitor (212). The networked computer system (200) is connected to a local area network (LAN) or a wide area network via a network interface connection (not shown). Those skilled in the art will appreciate that these input and output means may take other forms. Further, those skilled in the art will appreciate that one or more elements of the aforementioned computer (200) may be located at a remote location and connected to the other elements over a network. Further, software instructions to perform embodiments of the invention may be stored on a computer readable medium such as a compact disc (CD), a diskette, a tape, a file, or any other computer readable storage device.

[0062] While the invention has been described with respect to a limited number of embodiments, those skilled in the art, having benefit of this disclosure, will appreciate that other embodiments can be devised which do not depart from the

scope of the invention as disclosed herein. Accordingly, the scope of the invention should be limited only by the attached claims.

CLAIMS

What is claimed is:

1. A method for analyzing a target system, comprising:
   obtaining a plurality of characteristics from the target system using a characteristics extractor, wherein the plurality of characteristics is associated with a characteristics model;
   storing each of the plurality of characteristics in a characteristics store using a tracking mechanism; and
   analyzing the target system by issuing a query to the characteristics store to obtain an analysis result, wherein the query uses tracking information associated with the tracking mechanism.

2. The method of claim 1, further comprising:
   obtaining the characteristics model;
   generating the characteristics extractor associated with the characteristics model; and
   generating a characteristics store application programming interface (API) associated with the characteristics model, wherein the characteristics extractor uses the characteristics store API to store each of the plurality of characteristics in the characteristics store.

3. The method of claim 1, wherein the tracking mechanism comprises storing a time-based attribute associated with each of the plurality of characteristics in the characteristics store.

4. The method of claim 3, wherein the time-based attribute comprises at least one selected from the group consisting of a timestamp and a target system version number.

5. The method of claim 1, wherein the characteristics store comprises a first sub-characteristics store and a second sub-characteristics store.

6. The method of claim 5, wherein the tracking mechanism comprises storing characteristics associated with a first timeframe in the first sub-characteristics store and storing characteristics associated with a second timeframe in the second sub-characteristics store.

7. The method of claim 5, wherein the tracking mechanism comprises storing characteristics associated with a first version of the target system in the first sub-characteristics store and storing characteristics associated with a second version of the target system in the second sub-characteristics store.

8. The method of claim 1, wherein the tracking information comprises one selected from the group consisting of a time-related attribute and a sub-characteristics store name.

9. The method of claim 1, wherein the at least one query is defined using a pattern query language and wherein the pattern query language includes functionality to search for at least one pattern in the target system.

10. A system, comprising:

a characteristics model defining at least one artifact and a plurality of characteristics associated with the at least one artifact;

a target system comprising at least one of the plurality of characteristics defined in the characteristics model;

at least one characteristics extractor configured to obtain at least one of the plurality of characteristics from the target system;

a characteristics store configured to store the at least one of the plurality of characteristics obtained from the target system using a tracking mechanism; and

a query engine configured to analyze the target system by issuing a query to the characteristics store and configured to obtain an analysis result in response to the at least one query, wherein the query uses tracking information associated with the tracking mechanism.

11. The system of claim 10, wherein the tracking mechanism comprises storing a time-based attribute associated with each of the plurality of characteristics in the characteristics store.

12. The system of claim 11, wherein the time-based attribute comprises at least one selected from the group consisting of a timestamp and a target system version number.

13. The system of claim 10, wherein the characteristics store comprises a first sub-characteristics store and a second sub-characteristics store.

14. The system of claim 13, wherein the tracking mechanism comprises storing characteristics associated with a first timeframe in the first sub-characteristics store and storing characteristics associated with a second timeframe in the second sub-characteristics store.

15. The system of claim 13, wherein the tracking mechanism comprises storing characteristics associated with a first version of the target system in the first sub-characteristics store and storing characteristics associated with a second version of the target system in the second sub-characteristics store.

16. The system of claim 10, wherein the tracking information comprises one selected from the group consisting of a time-related attribute and a sub-characteristics store name.

17. The system of claim 10, wherein the at least one query is defined using a pattern query language, wherein the pattern query language includes functionality to search for at least one pattern in the target system.

18. The system of claim 10, wherein the characteristics model is a domain-specific model.

19. A computer readable medium comprising software instructions for analyzing a target system, comprising software instructions to:

obtain a plurality of characteristics from the target system using a characteristics extractor, wherein the plurality of characteristics is associated with a characteristics model;

store each of the plurality of characteristics in a characteristics store using a tracking mechanism; and

analyze the target system by issuing a query to the characteristics store to obtain an analysis result, wherein the query uses tracking information associated with the tracking mechanism.

20. The computer readable medium of claim 19, wherein the tracking mechanism comprises storing a time-based attribute associated with each of the plurality of characteristics in the characteristics store.

21. The computer readable medium of claim 20, wherein the time-based attribute comprises at least one selected from the group consisting of a timestamp and a target system version number.

22. The computer readable medium of claim 19, wherein the characteristics store comprises a first sub-characteristics store and a second sub-characteristics store.

23. The computer readable medium of claim 22, wherein the tracking mechanism comprises storing characteristics associated with a first timeframe in the first sub-

characteristics store and storing characteristics associated with a second timeframe in the second sub-characteristics store.

24. The computer readable medium of claim 22, wherein the tracking mechanism comprises storing characteristics associated with a first version of the target system in the first sub-characteristics store and storing characteristics associated with a second version of the target system in the second sub-characteristics store.

25. The computer readable medium of claim 19, wherein the tracking information comprises one selected from the group consisting of a time-related attribute and a sub-characteristics store name.

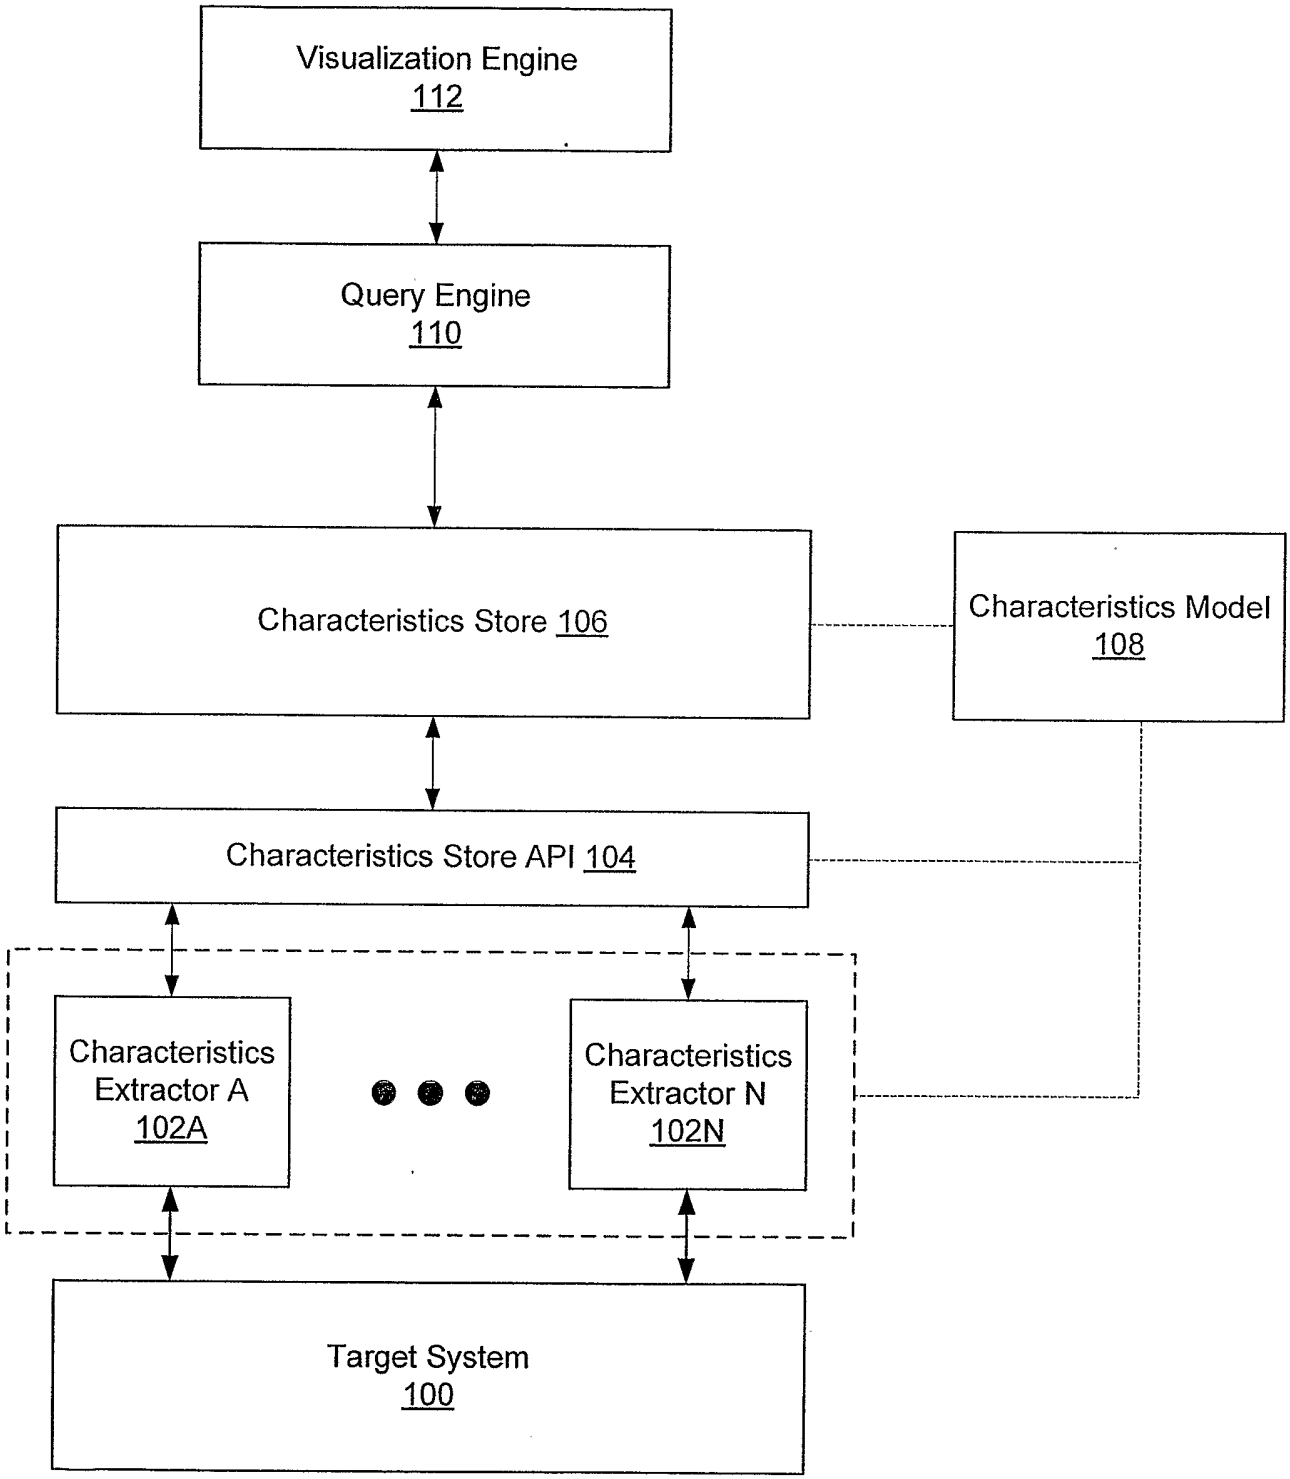

Visualization Engine
112
Query Engine
110
Characteristics Store 106
Characteristics Model
108
Characteristics Store API 104
Characteristics
Extractor A
102A
Characteristics
Extractor N
102N
Target System
100

FIG. 1

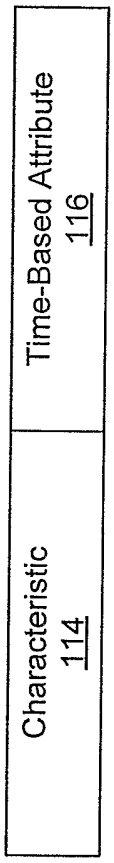

Characteristic
114
Time-Based Attribute
116

FIG. 2

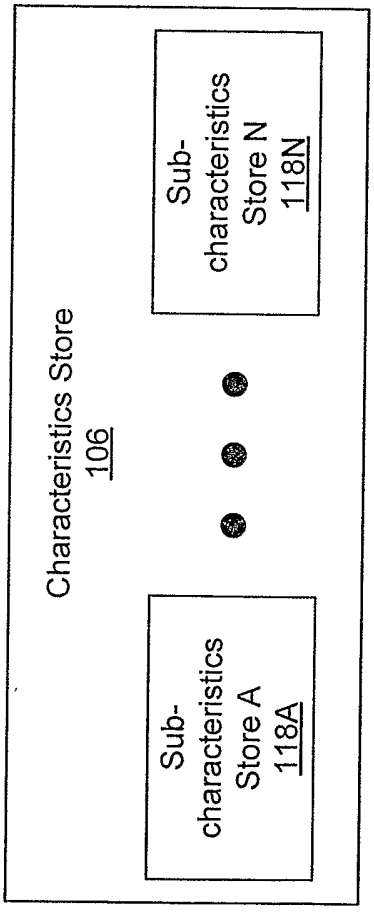

Characteristics Store
106
Sub-characteristics Store A
118A
Sub-characteristics Store N
118N

FIG. 3

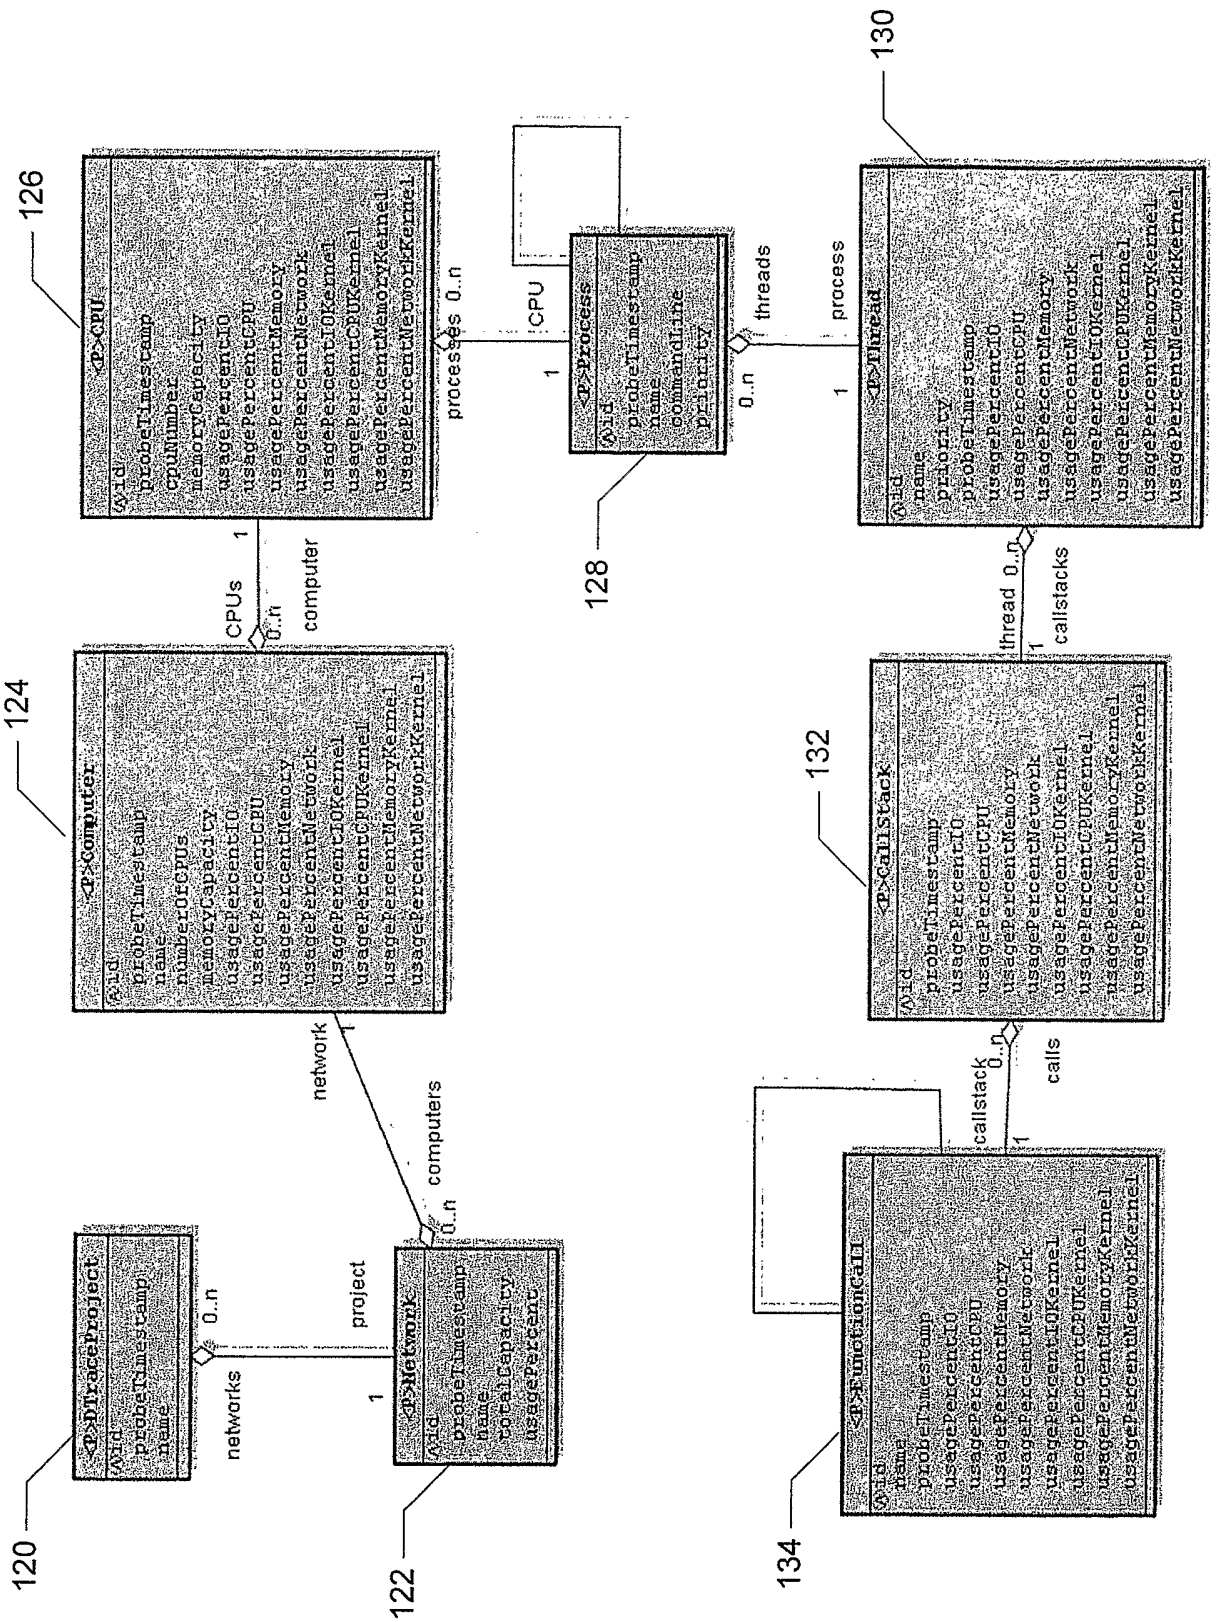

120
122
124
126
128
130
132
134
networks
project
network
computers
CPUs
computer
processes
CPU
threads
process
thread
callstacks
callstack
calls

FIG. 4

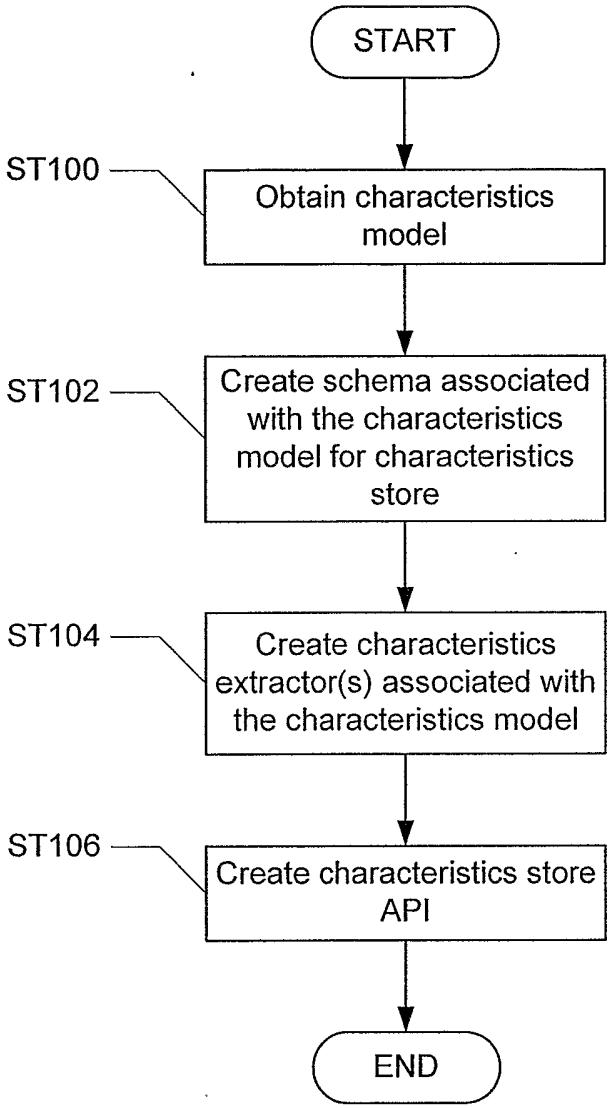

START
ST100
Obtain characteristics model
ST102
Create schema associated with the characteristics model for characteristics store
ST104
Create characteristics extractor(s) associated with the characteristics model
ST106
Create characteristics store API
END

FIG. 5

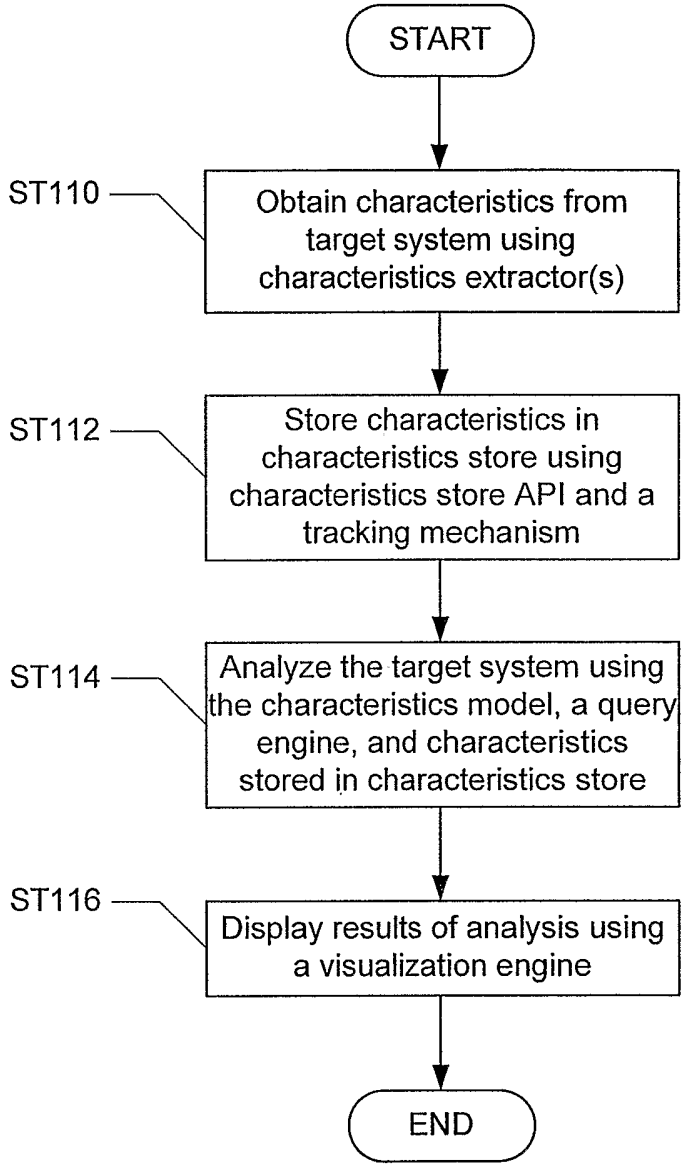

START
ST110
Obtain characteristics from target system using characteristics extractor(s)
ST112
Store characteristics in characteristics store using characteristics store API and a tracking mechanism
ST114
Analyze the target system using the characteristics model, a query engine, and characteristics stored in characteristics store
ST116
Display results of analysis using a visualization engine
END

FIG. 6

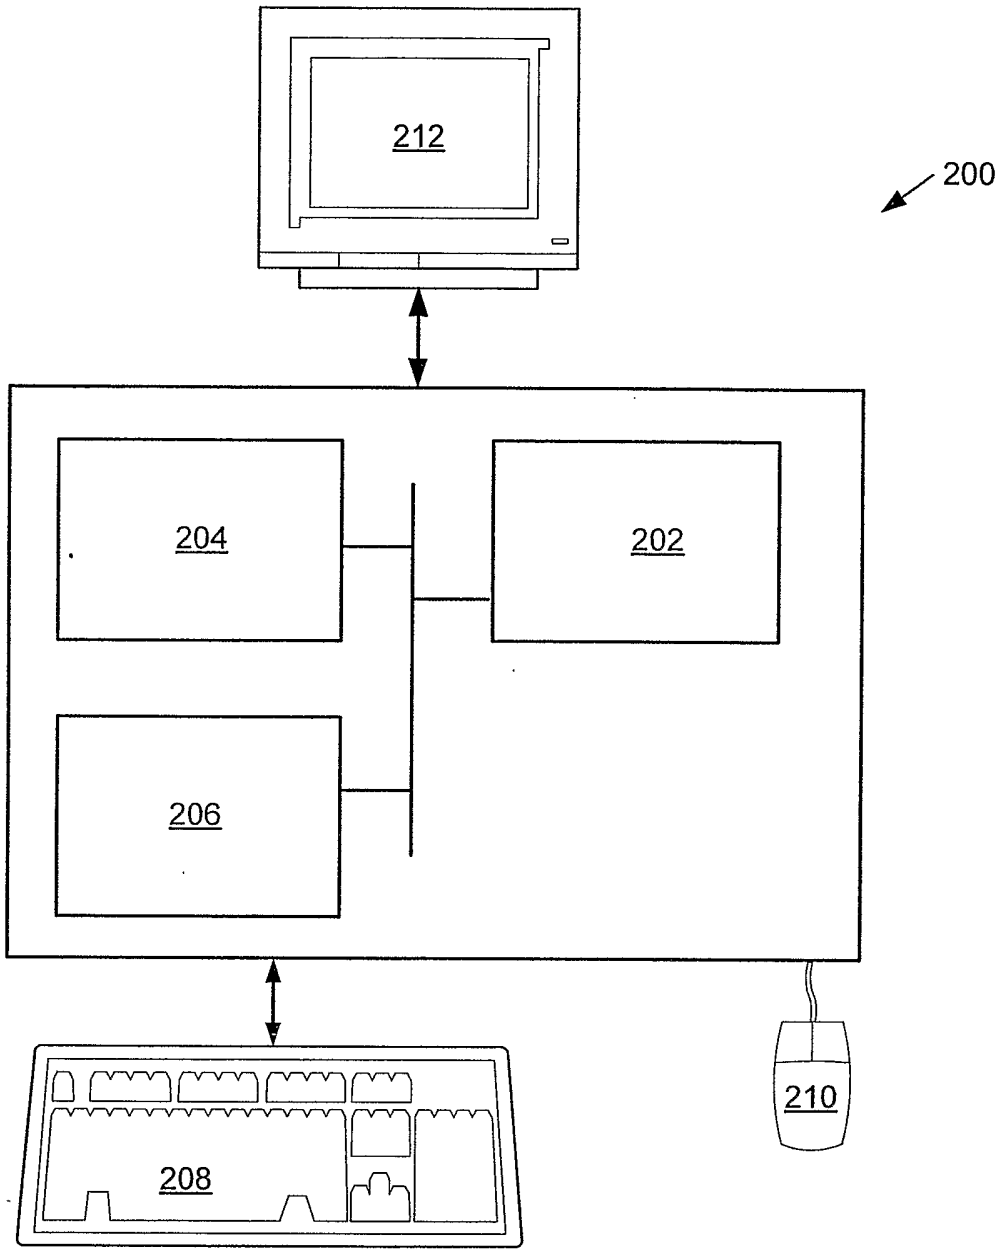


FIG. 7

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2005/018009

A. CLASSIFICATION OF SUBJECT MATTER
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, INSPEC, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 752 245 A (PARRISH ET AL) 12 May 1998 (1998-05-12)	1-4,8, 10-12, 16, 18-21,25
Y	abstract column 2, line 50 - line 67 column 7, line 52 - line 65 column 9, line 49 - column 10, line 13 column 10, line 35 - column 11, line 22 column 12, line 66 - column 13, line 4	5-7,9, 13-15, 17,22-24
Y	GB 2 383 152 A (* ORACLE CORPORATION) 18 June 2003 (2003-06-18) abstract	5-7, 13-15, 22-24

-/--

[X] Further documents are listed in the continuation of Box C.

[X] 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	Date of mailing of the international search report
23 January 2006	15/02/2006

Name and mailing address of the ISA/	Authorized officer
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	Renault, S

Form PCT/ISA/210 (second sheet) (April 2005)

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2005/018009

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 5 664 173 A (FAST ET AL) 2 September 1997 (1997-09-02) abstract	9,17
A	ATKINS D L: "Version sensitive editing: change history as a programming tool" SYSTEM CONFIGURATION MANAGEMENT. ECOOP'98 SCM-8 SYMPOSIUM. PROCEEDINGS 1998, 1998, pages 146-157, XP002363863 Berlin, Germany	

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2005/018009

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5752245 A	12-05-1998	CA 2205185 A1	13-06-1996
		DE 69503065 D1	23-07-1998
		DE 69503065 T2	21-01-1999
		EP 0786109 A2	30-07-1997
		JP 10512068 T	17-11-1998
		WO 9618157 A2	13-06-1996
GB 2383152 A	18-06-2003	US 2005262165 A1	24-11-2005
		US 2003115223 A1	19-06-2003
US 5664173 A	02-09-1997	NONE	