

(19) World Intellectual Property  
Organization  
International Bureau



(43) International Publication Date  
6 May 2005 (06.05.2005)

PCT

(10) International Publication Number  
**WO 2005/041061 A1**

(51) International Patent Classification<sup>7</sup>: **G06F 17/30**

[—/US]; 385 1st Avenue, Apt. 7G, New York, New York 10010 (US).

(21) International Application Number:

PCT/EP2004/052194

(74) **Agent: BURT, Roger, James;** IBM United Kingdom Limited, Intellectual Property Law, Hursley Park, Winchester, Hampshire SO21 2J (GB).

(22) International Filing Date:

15 September 2004 (15.09.2004)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

10/689,091

20 October 2003 (20.10.2003) US

(71) **Applicant (for all designated States except US): INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, New York 10504 (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, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(71) **Applicant (for MG only): IBM UNITED KINGDOM LIMITED** [GB/GB]; PO Box 41, North Harbour, Portsmouth, Hampshire PO6 3AU (GB).

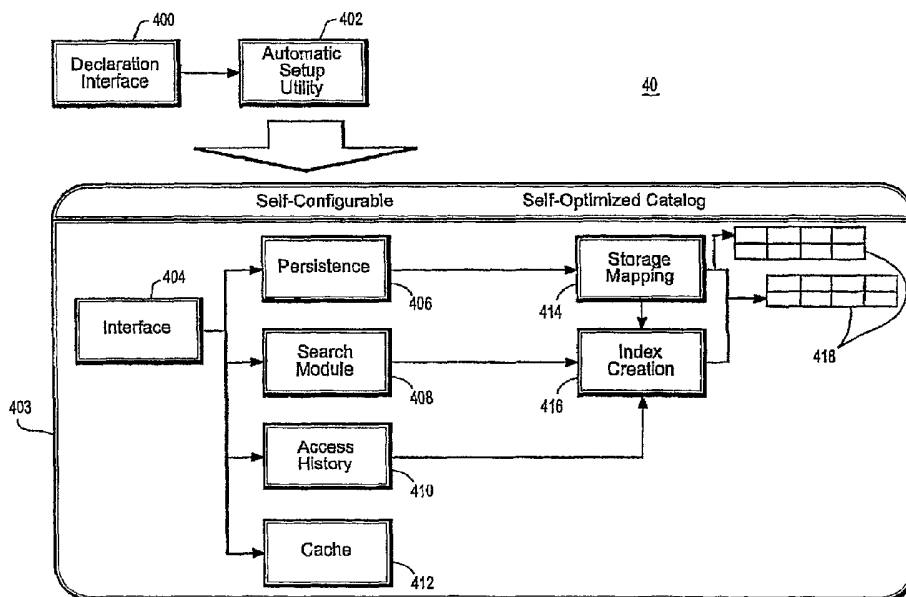
(84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT,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

(72) **Inventor; and**

(75) **Inventor/Applicant (for US only): CHANG, Yuan-Chi**

[Continued on next page]

(54) **Title:** SYSTEMS AND METHODS FOR PROVIDING AUTONOMOUS PERSISTENT STORAGE SYSTEMS



(57) **Abstract:** Systems, methods and services for generating autonomous persistent storage systems that are self-configurable and self-managing, based on user-submitted entity definitions. For example, systems and methods are provided for automatically creating and updating persistent storage structures based on entity definitions, automatically populating persistent storage space with instance data of defined entities, automatically generating and adapting methods for accessing instance data in persistent storage, searching instance data and automatically optimizing search methods for instance data, and automatically creating and managing a cache of frequently accessed instance data.



SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ,  
GW, ML, MR, NE, SN, TD, TG).

*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.*

**Published:**

— *with international search report*

## **Description**

# **SYSTEMS AND METHODS FOR PROVIDING AUTONOMOUS PERSISTENT STORAGE SYSTEMS**

### **Technical Field**

- [001] The present invention broadly relates to systems and methods for generating autonomous persistent storage systems that are self-configurable and self-managing. For example, the invention relates to systems and methods for providing autonomous creation, management, maintenance, and optimization of electronic catalogs for implementation with database systems and other information systems.

### **Background Art**

- [002] Typically, enterprise software systems such as Human Resource management (HR), Enterprise Resource Planning (ERP), Customer Relationship Management (CRM), Product Lifecycle Management (PLM), and Electronic Commerce (EC), include business management software, such as catalog functions, for managing a database. These enterprise software systems are typically installed and used at companies in many different industries, each of which having unique requirements. These heterogeneous requirements for enterprise systems creates a dilemma from a product development viewpoint, since it would be too bulky and complex to design business management software to accommodate all the various requirements for enterprise systems. On the other hand, a product should have a degree of flexibility to cover a minimal set of common requirements. In the end, a balance is made between the actual features that are shipped with the product releases and the anticipated industry/company-specific customization. One rule of thumb is that for large enterprise software, the cost to customize is three to five times larger than the cost to license a product. This high customization expense has prevented small and medium-sized business from aggressively adopting modern business management software.
- [003] The catalog function is one of the software components of an enterprise application that is most commonly customized to meet different needs. For example, with commerce applications in the retail industry, a media chain store may sell CDs and video, wherein the most important catalog data content includes singers, song lists, directors, producers, etc. In contrast, an office supply company may sell pens, pencils, folders, and other common office supplies, wherein the common catalog content concentrates on color, material and manufacturer, for example. The access and storage of information regarding singers, song lists, color and material is quite different. It is thus conceivably difficult to design a catalog to efficiently handle content related to CDs, video, pens, pencils, folders, as well as the millions of other products in retail and other

industries.

- [004] One reason for the constant struggle between catalog uniformity and product heterogeneity is due to the rigidity of traditional product catalog design. This is evidenced from a number of publications on electronic catalog as well as from the present inventor's experience with electronic commerce server development. The variations of schema design described, for example, in the article by S. Danish, "Building database-driven electronic catalogs," SIGMOD Record, Vol. 27, No. 4, December 1998, and in A. Jhingran, "Anatomy of a real E-commerce system," ACM SIGMOD Proceedings, May 2000) are common choices. But, these techniques and other conventional systems do not adapt well to product heterogeneity.

### **Disclosure of Invention**

- [005] The present invention provides a method as claimed in claim 1.
- [006] Exemplary embodiments of the present invention include systems and methods for generating autonomous persistent storage systems that are self-configurable and self-managing. Exemplary embodiments of the present invention include systems and methods for providing autonomous creation, management, maintenance, and optimization of persistent storage structures, such as directories of objects or electronic catalogs, for implementation with database systems and other information systems.
- [007] For example, exemplary embodiments of the invention include systems, methods, and services for automatically creating persistent storage structures and various persisted data management functions, based on user-submitted entity definitions. The entity definition may comprise, for example, a declaration of an object, one or more properties of the object, and a data type for each property. In other exemplary embodiments of the invention, the persistent storage structure may be, e.g., a database table or a file directory.
- [008] Exemplary systems and methods according to the invention provide mechanisms for automating the catalog lifecycle providing full adaptability to, e.g., heterogeneous product catalog needs, while eliminating the need for human developers to design persistent storage schemes and write software programs for access.
- [009] These and other exemplary embodiments, aspects, features, and advantages of the present invention will become apparent from the following detailed description of exemplary embodiments, which is to be read in connection with the accompanying drawings.

### **Brief Description of the Drawings**

- [010] FIG. 1 is a schematic diagram illustrating an operating system that can be used for implementing autonomous persistent storage systems and methods according to exemplary embodiments of the present invention.

- [011] FIG. 2 is an exemplary diagram illustrating functional relationships among entities that may access and interact with an autonomous persistent storage system according to an exemplary embodiment of the present invention.
- [012] FIG. 3 is an exemplary diagram illustrating an entity-property declaration that may be used for generating an autonomous catalog system according to an exemplary embodiment of the invention.
- [013] FIG. 4 is a schematic diagram of an autonomous electronic catalog system for a database according to an exemplary embodiment of the invention, which illustrates an exemplary embodiment of an autonomous persistent storage system according to the invention.
- [014] FIG. 5 is a flow diagram illustrating a method for parsing a new entity declaration into name and data type structures for use in generating an autonomous catalog system, according to an exemplary embodiment of the invention.
- [015] FIG. 6 is a flow diagram illustrating a method for automatically creating a persistent storage structure in the form of a database table according to an exemplary embodiment of the invention.
- [016] FIG. 7 is a flow diagram illustrating a method for automatically generating a software class that provides methods for accessing, searching and deleting entity instances, according to an exemplary embodiment of the invention.
- [017] FIG. 8 is a flow diagram illustrating a method for automatically populating instance data into a database table, according to an exemplary embodiment of the invention.
- [018] FIG. 9 is a flow diagram illustrating a method for automatically creating an index for object instance data based on access history, according to an exemplary embodiment of the invention.
- [019] FIG. 10 is a flow diagram illustrating a method for implementing a cache to optimize instance access performance, according to an exemplary embodiment of the invention.

### **Best Mode for Carrying Out the Invention**

- [020] Exemplary embodiments of the present invention include systems and methods for generating autonomous persistent storage systems, which are self-configurable, self-managing and self-optimizable. For example, exemplary embodiments of the present invention include systems and methods for providing autonomous creation, management, maintenance, and optimization of persistent storage structures, such as directories of objects or electronic catalogs, for implementation with database systems and other information systems.
- [021] In general, exemplary embodiments of the invention include methods for automatically generating persistent storage structures, e.g., electronic catalogs, using an

entity description of an entity and associated properties to be stored in a catalog, persistent storage of entity instances that comply with the entity description, as well as methods for automatically, configuring and managing the persistent storage and synthesizing an interface for accessing persisted entity instances.

[022] It is to be understood that the term "database" as used herein is to be broadly construed to include any persistent storage system, such as, for example, relational databases (e.g., IBM DB2 and Oracle 9), file systems (e.g., NTFS), disk management systems (e.g., RAID), etc. For illustrative purposes, exemplary embodiments of the present invention are described with reference to relational database systems, but those of ordinary skill in the art can readily envision other types of persistent storage systems and methods that can be used in lieu of relational databases.

[023] Furthermore, the term "catalog" or "electronic catalog" as used herein should be broadly construed to refer to any directory of entities characterized by properties. For example, a "catalog" may be structured as a database table. It is to be further understood that the term "entity" as used herein is to be broadly construed as referring to any "thing" or "object" that is to be persistently stored. Further, the term "property" as used herein is intended to broadly refer to attributes of an entity. By way of example, an *entity* named "soda" may be defined having *properties* (attributes) such as *name, flavor, size, manufacturer, nutrition fact, UPC* (Universal Product Code), etc. The term "instance" as used herein broadly refers to a physical materialization of an entity. Using the "soda" entity as an example, an *instance* of "soda" can be, e.g., Diet Coke, manufactured by the Coca Cola Company.

[024] It is to be appreciated that exemplary embodiments of the present invention as described herein can be used to implement, for example, catalog functions in commerce servers, human resource management, enterprise resource planning, procurement systems, supply chain systems, customer relationship management, product lifecycle management, as well as other enterprise software applications. For purposes of illustration and description, exemplary embodiments of the invention will be described herein based on a Java (Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries or both) implementation and relational database, although nothing herein shall be construed as limiting the scope of the invention with regard to, e.g., the programming language or persistent storage system.

[025] It is to be understood that the exemplary systems and methods described herein may be implemented in various forms of hardware, software, firmware, special purpose processors, or a combination thereof. In particular, the present invention can be implemented as an application comprising program instructions that are tangibly embodied on one or more program storage devices (e.g., hard disk, magnetic floppy

disk, RAM, ROM, CD ROM, etc.) and executable by any device or machine comprising suitable architecture.

[026] It is to be further understood that, because some of the constituent system components and process steps depicted in the accompanying Figures are preferably implemented in software, the connections between system modules (or the logic flow of method steps) may differ depending upon the manner in which the present invention is programmed. Given the teachings herein, one of ordinary skill in the related art will be able to contemplate these and similar implementations or configurations of the present invention.

[027] For example, FIG. 1 is a schematic diagram illustrating an operating system that can be used for implementing an autonomous persistent storage system (e.g., autonomous catalog system) according to an exemplary embodiment of the present invention. The system (10) comprises a user interface (100), CPU (central processing unit) (102), a cache/RAM memory (104) and persistent storage memory (106). The user interface (100), which comprises a display and keyboard, enables a user to interface with an autonomous catalog system which executes in the CPU (102). It is to be understood that the autonomous catalog system may interface with other applications or services using suitable APIs and protocols for computer networking communications, for example. An autonomous catalog system according to an exemplary embodiment of the present invention executes in the CPU (102) and uses the main memory (104) to temporarily store its executable code and data. Furthermore, the disk (106) is used for storing programs and data that require persistence. It is to be understood that the system (10) depicted in Fig. 1 is merely one exemplary embodiment of an operating environment in which autonomous catalog systems and methods according to the present invention can be implemented and executed, and that other suitable operating environments for use with the present invention are can be readily envisioned by one of ordinary skill in the art.

[028] Referring now to Fig. 2, an exemplary diagram illustrates functional relationships among entities that may access and interact with an autonomous persistent storage system according to an exemplary embodiment of the present invention. Fig. 2 depicts an abstract view of a catalog (214) comprising a directory of entity *instances*. The *instances* may or may not belong to the same *entity*. In accordance with one exemplary embodiment of the invention, an autonomous catalog system supports various access methods including, for example, store (208), retrieve (210), search (212) and remove (216). The store (208) method is used to insert new instances into the catalog (214). The retrieve (210) method is used to obtain a specified instance from the catalog (214) by identifiers. The search (212) method is used to find instances whose property values match specified conditions. The remove (216) method can be optionally implemented

to delete specified instances from the catalog (214) by identifiers. The various methods store (208), retrieve (210), search (212) and/or remove (216) may be used various entities, including, for example, human operators (200), computer hardware (202), computer software (204), web services (206), etc., depending on the implementation or purpose of the catalog (214) in a larger enterprise software.

[029] Referring now to FIG. 3, an exemplary diagram illustrates an entity description (e.g., entity-property declaration), which may be used for generating an autonomous catalog system, according to an exemplary embodiment of the invention. In the illustrative embodiment of Fig. 3, an entity-property declaration (30) comprises an *entity* (300) having a unique name (302) for identification. The *entity* (300) comprises a plurality of *properties* (304-1, 304-2, ..., 304-N). Each *property* (304-1, 304-2, ..., 304-N) is identified by a corresponding *property name* (306-1, 306-2, ... 306-N) and is assigned a corresponding *data type* (307-1, 307-2, ..., 307-N). The *data type* field is used for estimating the amount of storage space for the *property* value. The data types may include, for example, integer, floating point, character strings, byte arrays and other data types known by those of ordinary skill in the art. In one exemplary embodiment of the present invention, new data types can be defined as needed, depending on the application.

[030] It is to be appreciated that in one exemplary embodiment of the invention, the entity-property declaration (30) can be generated using application-dependent methods that are compatible with the programming models of the enterprise software that implements an autonomous catalog system according to the invention. In other exemplary embodiments, the entity-property declaration (30) can be generated using application-independent languages, such as XML, for example, or other suitable extensible languages. It is to be understood that Fig. 3 is merely one illustrative embodiment of an entity description including information that is preferable for implementing an autonomous catalog according to an embodiment of the invention, and that additional information may be used for generating an entity-property declaration depending on the application.

[031] Fig. 4 depicts a high-level schematic diagram of an autonomous catalog system according to an exemplary embodiment of the present invention. It is to be understood that Fig. 4 is merely one exemplary embodiment of an autonomous persistent storage system in the form of an electronic catalog system for a relational database. However, as noted above, one of ordinary skill in the art can readily envision other autonomous persistent storage systems that can be implemented in accordance with the teachings of the present invention.

[032] Referring now to Fig. 4, the autonomous catalog system (40) comprises a declaration interface (400) and automatic setup utility (402), which process an entity



description to generate an autonomous (self-configurable, self-managed, self-optimized) electronic catalog (403). More specifically, in one exemplary embodiment, self-configuration of the autonomous catalog (403) begins with the declaration interface (400) receiving an entity definition and then interpreting and parsing the entity definition to form logical associations. For example, the declaration interface (400) can parse an entity definition to determine the structure and properties of the data and generate a hierarchically structured tree as depicted in FIG. 3, for example. The declaration interface (400) passes the parsing results (logical associations) to the automatic setup utility (402), which commences a configuration process to generate the autonomous catalog (403) having a plurality of modules.

[033] For example, the autonomous catalog (403) comprises an interface module (404) having access programs that are synthesized to enable access to persisted entities. For example, the interface module (404) comprises methods to *get* and *set property* values, as well as *searchby* property values of persisted entities. A persistence module (406) and storage mapping module (414) function jointly to provide a reliable utility for storing catalog entity instances in database tables (418). A search module (408) provides a method for searching catalog entity instances in the database tables (418). An index creation module (416) provides a method for indexing one or more catalog entity instances, and can be invoked by the search module (408) to improve search performance. An access history module (410) provides a mechanism for controlling the index creation module (416) to create indexes for frequently searched properties. A cache module (412) provides a mechanism for storing recently accessed instances in database tables (418).

[034] In one exemplary embodiment, the system (40) of Fig. 4 comprises modules (400), (402) and (403) comprising modules (404), (406), (414), (408) and (418). In another exemplary embodiment, modules (410), (412), and (416), for example, can be implemented to enhance access and retrieval performance.

[035] It is appreciated that an autonomous persistent storage system, such as depicted in Fig. 4, can be employed for various applications. For instance, such system can be implemented as a component of an enterprise application for providing database management. In addition, an autonomous persistent storage system can be implemented with web/application servers for building e-commerce applications. In addition, an autonomous persistent storage system may be implemented as a service that is accessible on a remote server, wherein a consumer can pay service fees to a service provider based on some fee structure or service level agreement for providing, e.g., secured database management. Those of ordinary skill in the art can readily envision other applications for an autonomous persistent storage system according to the invention.

- [036] Details regarding exemplary modes of operation, functions and architectures of the autonomous catalog system (40) modules will now be described in further detail with reference to the exemplary embodiments depicted in Figs. 5-10, for example. In addition, for illustrative purposes, the exemplary systems and methods depicted in Figs. 5-10 will be described with reference to an exemplary entity-property declaration for a declared *entity* Soda, having *properties* Brand, Size, and UPC. In addition, it is assumed that the *property* Brand has a “string” data type, the *property* Size has an “integer” data type and the *property* UPC has a “string” data type.
- [037] Fig. 5 is a flow diagram illustrating a method for parsing a new entity definition into name and data type structures for use in generating an autonomous catalog system according to an exemplary embodiment of the invention. For example, Fig. 5 illustrates a system and method for implementing the declaration interface module (400) (Fig. 4) according to an exemplary embodiment of the invention.
- [038] Referring now to Fig. 5, initially, a new entity declaration (500) is received for processing. As noted above, the entity declaration (500) may be described using an application-dependent language or application-independent language. The entity declaration (500) is parsed to determine the “name” of the declared entity (502), and the parsed entity name (504) is stored for subsequent access in downstream processing. In the above example, the entity name Soda will be stored. Further, the declaration is parsed to obtain the property names (506) of the declared entity. For each property name, the declaration is further parsed to obtain the corresponding data type, and the parsed property names and corresponding data types (510) are stored for subsequent access in downstream processing. In the above example, the entity properties Brand, Size, and UPC (and corresponding data types) will be parsed and stored.
- [039] Fig. 6 is a flow diagram illustrating a method for creating persistent storage structure according to an exemplary embodiment of the invention. For example, FIG. 6 illustrates a system and method for implementing the persistence (406) and storage mapping (414) modules (Fig. 4) for creating a database table for a relational database according to exemplary embodiments of the invention. For purposes of discussion, with no loss of generality, it is assumed that the persistent storage is a relational database that accepts SQL (Structured Query Language) queries/requests.
- [040] Referring now to Fig. 6, initially, a database connection is opened (600). A SQL statement for creating a new database table is automatically generated using, for example, the parsed entity name (504) and the entity property names and associated data types (510). More specifically, a create table SQL statement is generated by defining a new database table with the entity name (602). In one exemplary embodiment using SQL and a relational database, and using the entity name Soda in the above example, a create table SQL statement begins with “CREATE TABLE

SODA". In other exemplary embodiments of the invention, the entity name (504) can be transformed or altered to comply with certain requirements of the database management system, as is understood by those of ordinary skill in the art (a detailed discussion of which is beyond the scope of the present invention).

[041] Then, generation of the create table statement continues by defining a new column for the database table for each property name and corresponding data type (604). More specifically, a table column is created for each property based on its data type and the property names are used as column names. For example, since the entity Soda comprises the properties Brand, Size and UPC, the create table SQL statement is further expanded as follows:

[042] "(BRAND CHARACTER (10) NOT NULL, SIZE BIGINT NOT NULL, UPC CHARACTER (10) NOT NULL)".

[043] In the exemplary SQL statement, the Brand column has enough space for ten characters. The Size column can store an integer. The UPC column also has enough space for ten characters. Finally, the exemplary SQL statement is appended with database user space clause. The complete table creation statement is as follows:

[044] "CREATE TABLE SODA (BRAND CHARACTER (10) NOT NULL, SIZE BIGINT NOT NULL, UPC CHARACTER (10) NOT NULL) IN "USERSPACE";"

[045] Once the table creation statement is generated, the statement is issued to the database to create a table (606) having a structure as defined in the statement. In the above example, it is assumed that a database table SODA is created.

[046] FIG. 7 is a flow diagram illustrating a method for automatically generating a software class that provides methods for accessing, searching and deleting entity instances, according to an exemplary embodiment of the invention. For example, FIG. 7 illustrates a system and method for implementing the interface module (404) (Fig. 4) for the autonomous catalog (403) according to an exemplary embodiment of the invention. For purposes of discussion, without loss of generality, it is assumed that the Java programming language is employed for the autonomous catalog system of Fig. 4.

[047] Referring now to Fig. 7, initially, a new software (Java) class is created (700), which is named after the stored entity name (504) (i.e., Soda using the above example). One or more members of the new class will then be declared with the names and data types of entity properties (510). More specifically, for each property, a new field is created with the property name and associated data type (702). By way of example, using the above Soda example, an exemplary Java skeleton of the class is as follows:

```
[048] public class Soda {  
[049]     public String Brand;  
[050]     public int Size;  
[051]     public String UPC;
```

[052] }

[053] Next, for each field (property), a pair of get() and put() methods are created (706). In addition, a searchBy() method is created for each property (708). Further, a deleteSelf() method will be created (704) so that the class can destroy its own instance. By way of example using the above Soda example, the exemplary Java skeleton of the class is expanded as follows:

```
[054] public class Soda {  
[055]     public String Brand;  
[056]     public int Size;  
[057]     public String UPC;  
[058]     public String getBrand() {  
[059]         return Brand; }  
[060]     public int getSize() {  
[061]         return Size; }  
[062]     public String getUPC() {  
[063]         return UPC; }  
[064]     public void setBrand(String brand) {  
[065]         Brand = brand; }  
[066]     public void setSize(int size) {  
[067]         Size = size; }  
[068]     public void setUPC(String uPC) {  
[069]         UPC = uPC; }  
[070]     public Soda [] searchByBrand(String brand) {  
[071]         return null; }  
[072]     public Soda [] searchBySize(int size) {  
[073]         return null; }  
[074]     public Soda [] searchByUPC(String uPC) {  
[075]         return null; }  
[076]     public void deleteSelf() {  
[077]     }  
[078] }
```

[079] Further, a populate() method is created to insert property values into the SODA table (712). The exemplary embodiment of Fig. 8 illustrate an implementation of the populate() method. In addition, the deleteSelf() method is populated with a storage delete command (710), e.g., an SQL statement to delete the row representing the instance from the SODA table (710). A procedure to generate the SQL update is similar to FIG. 8 and thus will not be repeated here. For completeness, an exemplary deleteSelf() statement is illustrated for reference:

[080]       “DELETE FROM SODA WHERE BRAND = this.Brand AND SIZE = this.Size  
AND UPC = this.UPC”

[081]       The exemplary Java class skeleton, which has the populate() and deleteSelf()  
method populated, is as follows:

```
[082]       public class Soda {  
[083]           public String Brand;  
[084]           public int Size;  
[085]           public String UPC;  
[086]           public String getBrand() {  
[087]               return Brand; }  
[088]           public int getSize() {  
[089]               return Size; }  
[090]           public String getUPC() {  
[091]               return UPC; }  
[092]           public void setBrand(String brand) {  
[093]               Brand = brand; }  
[094]           public void setSize(int size) {  
[095]               Size = size; }  
[096]           public void setUPC(String uPC) {  
[097]               UPC = uPC; }  
[098]           public Soda [] searchByBrand(String brand) {  
[099]               return null; }  
[100]           public Soda [] searchBySize(int size) {  
[101]               return null; }  
[102]           public Soda [] searchByUPC(String uPC) {  
[103]               return null; }  
[104]           public void deleteSelf() {  
[105]               execute the SQL statement “DELETE FROM SODA WHERE BRAND =  
this.Brand AND SIZE = this.Size AND UPC = this.UPC”;  
[106]               }  
[107]           public void populate() {  
[108]               execute the SQL statement “INSERT INTO SODA (BRAND, SIZE, UPC)  
VALUES (this.Brand, this.Size, this.UPC)”;  
[109]               }  
[110]           }
```

[111]       Further, a populate() method is created to insert property values into the SODA  
table (712). The exemplary embodiment of Fig. 8 illustrate an implementation of the  
populate() method. In addition, the deleteSelf() method is populated with a storage

delete command (710), e.g., an SQL statement to delete the row representing the instance from the SODA table (710). A procedure to generate the SQL update is similar to FIG. 8 and thus will not be repeated here. For completeness, an exemplary deleteSelf() statement is illustrated for reference:

[112]       “DELETE FROM SODA WHERE BRAND = this.Brand AND SIZE = this.Size  
AND UPC = this.UPC”

[113]       The exemplary Java class skeleton, which has the populate() and deleteSelf() method populated, is as follows:

```
[114]       public class Soda {  
[115]           public String Brand;  
[116]           public int Size;  
[117]           public String UPC;  
[118]           public String getBrand() {  
[119]               return Brand; }  
[120]           public int getSize() {  
[121]               return Size; }  
[122]           public String getUPC() {  
[123]               return UPC; }  
[124]           public void setBrand(String brand) {  
[125]               Brand = brand; }  
[126]           public void setSize(int size) {  
[127]               Size = size; }  
[128]           public void setUPC(String uPC) {  
[129]               UPC = uPC; }  
[130]           public Soda [] searchByBrand(String brand) {  
[131]               return null; }  
[132]           public Soda [] searchBySize(int size) {  
[133]               return null; }  
[134]           public Soda [] searchByUPC(String uPC) {  
[135]               return null; }  
[136]           public void deleteSelf() {  
[137]               execute the SQL statement “DELETE FROM SODA WHERE BRAND =  
              this.Brand AND SIZE = this.Size AND UPC = this.UPC”;  
[138]               }  
[139]           public void populate() {  
[140]               execute the SQL statement “INSERT INTO SODA (BRAND, SIZE, UPC)  
              VALUES (this.Brand, this.Size, this.UPC)”;  
[141]               }
```

[142] }

[143] Furthermore, the searchBy() methods for the properties are populated with database query and retrieval statements (714). More specifically, by way of example, the searchBy() method looks for matches of Soda instances with specified conditions. For instance, the query searchByBrand("Diet Coke") requests all Soda instances whose Brand property is "Diet Coke". In one exemplary embodiment, the result that is returned comprises an array of Soda classes with the property values populated:

[144] public Soda [] searchByBrand(String brand) {

[145] execute the SQL statement "SELECT \* FROM SODA WHERE BRAND =  
"brand" ";

[146] For each instance match, instantiate a new Soda object;

[147] Set properties of the new Soda object with return values from the SQL statement;

[148] Return all the instantiated Soda objects as an array;

[149] }

[150] In another exemplary embodiment, the searchBy() method syntax can be extended using additional operators for fuzzy matching and number comparisons, for example, as is readily understood by those of ordinary skill in the art.

[151] The exemplary Java class skeleton, which has all the methods populated, is as follows:

[152] public class Soda {

[153] public String Brand;

[154] public int Size;

[155] public String UPC;

[156] public String getBrand() {

[157] return Brand; }

[158] public int getSize() {

[159] return Size; }

[160] public String getUPC() {

[161] return UPC; }

[162] public void setBrand(String brand) {

[163] Brand = brand; }

[164] public void setSize(int size) {

[165] Size = size; }

[166] public void setUPC(String uPC) {

[167] UPC = uPC; }

[168] public Soda [] searchByBrand(String brand) {

[169] execute the SQL statement "SELECT \* FROM SODA WHERE BRAND =  
"brand" ";

- [170] For each instance match, instantiate a new Soda object;
- [171] Set properties of the new Soda object with return values from the SQL statement;
- [172] Return all the instantiated Soda objects as an array;
- [173] }
- [174] public Soda [] searchBySize(int size) {
- [175] execute the SQL statement "SELECT \* FROM SODA WHERE SIZE = size";
- [176] For each instance match, instantiate a new Soda object;
- [177] Set properties of the new Soda object with return values from the SQL statement;
- [178] Return all the instantiated Soda objects as an array;
- [179] }
- [180] public Soda [] searchByUPC(String uPC) {
- [181] execute the SQL statement "SELECT \* FROM SODA WHERE UPC = "uPC" ";
- [182] For each instance match, instantiate a new Soda object;
- [183] Set properties of the new Soda object with return values from the SQL statement;
- [184] Return all the instantiated Soda objects as an array;
- [185] }
- [186] public void deleteSelf() {
- [187] execute the SQL statement "DELETE FROM SODA WHERE BRAND =  
this.Brand AND SIZE = this.Size AND UPC = this.UPC";
- [188] }
- [189] public void populate() {
- [190] execute the SQL statement "INSERT INTO SODA (BRAND, SIZE, UPC)  
VALUES (this.Brand, this.Size, this.UPC)";
- [191] }
- [192] }
- [193] In the above exemplary embodiment, the newly generated Soda class, which accesses the newly created Soda table, can now be called by other applications, for example, to populate, search and/or delete *instances* of Soda. Although the above exemplary systems and methods describe an autonomous method for creating a catalog for Sodas, those of ordinary skill in the art will readily appreciate that syntheses and creation of new classes and tables for additional declared entities can be performed using the same procedures.
- [194] FIG. 8 is a flow diagram illustrating a method for populating instance data into a database table, according to an exemplary embodiment of the invention. For example, Fig. 8 illustrates a system and method for implementing step (712) (Fig. 7). Continuing by way of example with the Soda example, it is assumed in Fig. 8 that the Soda database table has been created and that a new instance having catalog content is processed to populate such catalog content into the Soda table.



[195] Referring to Fig. 8, initially, a database connection is open (800). Then, creation of an SQL insert statement begins with "INSERT INTO" (802). A new entity instance declaration for Soda is received (804) and parsed to obtain the entity name (806). By way of example, the new entity instance for Soda may have the following properties and value: Brand = "Diet Coke", Size = 12 oz, and UPC= 129100. In the table insert statement, the database table is named with the entity name (808). The instance declaration is parsed to obtain the names and values for each property (810). The parsed property names and values are then used to generate the remainder of the insert statement (812). Using the above exemplary Soda instance, an exemplary SQL insert statement is as follows:

[196] "INSERT INTO SODA (BRAND, SIZE, UPC) VALUES ('Diet Coke', 12, '129100');"

[197] Alternatively, if the new instance has only Brand and Size information, the SQL insert statement can be altered to include column names as follows:

[198] "INSERT INTO SODA (BRAND, SIZE) VALUES ('Diet Coke', 12);"

[199] The insert statement is then issued to the database and the SQL operation "commit" is invoked to write the instance data to persistent memory (814).

[200] FIG. 9 is a flow diagram illustrating a method for automatically creating an index for object instance data based on access history, according to an exemplary embodiment of the invention. For example, FIG. 9 illustrates a system and method for implementing the access history (410) and index creation (416) modules (Fig. 4).

[201] Referring to now to Fig. 9, when a searchByFieldName() method is invoked for a given database table (900), a matched field counter for the requested field (property) name is incremented by one (902). In one exemplary embodiment, each field (table column) has a counter. If the counter number exceeds a preset threshold (affirmative determination in 904), an index will be created for the table field (property) (906). By way of the above Soda example for a relational database, a command for creating an index for the UPC column of the table Soda can be as follows:

[202] "CREATE INDEX I\_UPC ON SODA (UPC ASC) PCTFREE 10 MINPCTUSED 10;"

[203] On the other hand, if the count does not exceed the predefined threshold (negative determination in 904), no index will be created. It is to be appreciated that the exemplary systems and methods described in Fig. 9 can be implemented to enhance access performance for an autonomous persistent storage system according to the invention.

[204] FIG. 10 is a flow diagram illustrating a method for implementing a cache to optimize instance access performance, according to an exemplary embodiment of the invention. For example, FIG. 10 illustrates a system and method for implementing the

cache module (412) (Fig. 4). Initially, when a request for an entity object is received (1000), a cache memory will be searched for the requested object (1002). If the object is cached (affirmative determination in 1004), the object is retrieved from the cache (1006) and returned (1012). If the object is not cached (negative determination in 1002), a new entity object is instantiated (1008) by reading the object from the persisted memory to RAM, for example. The instantiated entity object is cached (1010) and returned (1012).

- [205] It is to be appreciated that the exemplary systems and methods described herein in accordance with the present invention may be efficiently and effectively implemented for providing persistent data management in a variety of enterprise applications such as human resource management (HR), enterprise resource management (ERP), customer relationship management (CRM), electronic commerce (EC), etc. Indeed, although such enterprise applications can implement a wide range of different catalog structures for various purposes, the present invention provides mechanisms for abstracting and automating catalog functions in the whole lifecycle, thereby enabling automatic and dynamic adaptation to product heterogeneity and, thus, minimizing the total cost of ownership.
- [206] For example, in one exemplary embodiment, the present invention provides mechanisms for automatically and dynamically creating new database tables and new access classes to adapt to product heterogeneity, without requiring manual customization. For example, an autonomous catalog system according to an embodiment of the invention can start with as few as two tables and grow automatically according to the needs of catalog users. An autonomous catalog system supports the basic definition of catalog, which is a directory of things characterized by their properties. By leaving the definition of a “thing” and its “properties” to catalog users, an autonomous catalog system according to the present invention provides the flexibility to support virtually all catalog requirements regardless of the industry domain.
- [207] In addition, an autonomous catalog system according to an embodiment of the invention enables an application developer to focus on the core functionality of an enterprise application, without having to know details of the schema design and performance optimization for the catalog system to be implemented with the enterprise application.

## Claims

- [001] A method for generating a persistent storage system, the method comprising the steps of: receiving as input an entity definition; automatically generating a persistent storage structure for a persistent storage medium using the entity definition; and automatically generating an interface for accessing the persistent storage medium.
- [002] The method of claim 1, wherein the entity definition comprises a declaration of an object, one or more properties of the object, and a data type for each property.
- [003] The method of claim 1, wherein the persistent storage structure comprises a database table.
- [004] The method of claim 1, wherein the persistent storage structure comprises a file directory.
- [005] The method of claim 1, wherein the persistent storage medium comprises a hard disk, a readable/writeable CD or a floppy disk.
- [006] The method of claim 1, wherein the method is implemented in a database system.
- [007] The method of claim 6, wherein the database system is a relational database.
- [008] The method of claim 1, wherein the step of automatically generating an interface for accessing the persistent storage medium comprises automatically creating methods for one of storing, retrieving, searching, and removing entity instance data in the persistent storage medium.
- [009] The method of claim 1, further comprising the step of automatically generating an index to persistent stored data.
- [010] The method of claim 1, wherein the step of automatically generating an index to persistent stored data comprises generating an index to persistent stored data that is frequently accessed as determined by a predefined indicator.
- [011] The method of claim 1, further comprising the step of automatically adapting the persistent storage structure or the access interface for a new entity definition.
- [012] The method of claim 1, further comprising the step of automatically optimizing the persistent storage system to improve search efficiency or storage scalability.
- [013] The method of claim 1, further comprising automatically creating a cache memory for storing an entity instance that is accessed from the persistent storage medium.
- [014] The method of claim 1, further comprising the steps of: receiving an entity instance declaration; and automatically populating the persistent storage structure with entity instance data.
- [015] A program storage device readable by a machine, tangibly embodying a program of instructions executable by the machine to perform method steps for generating

- a persistent storage system, the method comprising the steps of: receiving as input an entity definition; automatically generating a persistent storage structure for a persistent storage medium using the entity definition; and automatically generating an interface for accessing the persistent storage medium.
- [016] The program storage device of claim 15, wherein the entity definition comprises a declaration of an object, one or more properties of the object, and a data type for each property.
- [017] The program storage device of claim 15, wherein the persistent storage structure comprises a database table.
- [018] The program storage device of claim 15, wherein the persistent storage structure comprises a file directory.
- [019] The program storage device of claim 15, wherein the instructions for automatically generating an interface for accessing the persistent storage medium comprise instructions for automatically creating methods for one of storing, retrieving, searching, and removing entity instance data in the persistent storage medium.
- [020] The program storage device of claim 15, further comprising instructions for automatically generating an index to persistent stored data.
- [021] The program storage device of claim 20, wherein the instructions for automatically generating an index to persistent stored data comprise instructions for generating an index to persistent stored data that is frequently accessed as determined by a predefined indicator.
- [022] The program storage device of claim 15, further comprising instructions for automatically adapting the persistent storage structure or the access interface for a new entity definition.
- [023] The program storage device of claim 15, further comprising instructions for automatically optimizing the persistent storage system to improve search efficiency or storage scalability.
- [024] The program storage device of claim 15, further comprising instructions for automatically creating a cache memory for storing an entity instance that is accessed from the persistent storage medium.
- [025] The program storage device of claim 15, further comprising instructions for performing the steps of: receiving an entity instance declaration; and automatically populating the persistent storage structure with entity instance data.
- [026] A persistent storage system, comprising: an interface for receiving and processing an entity definition; and a utility module for automatically configuring an autonomous persistent storage system in accordance with the processed entity definition, wherein the autonomous persistent storage system is

automatically configured to comprise: a module for automatically generating a persistent storage structure for a persistent storage medium using the entity definition; and an access interface for invoking access methods for accessing the persistent storage medium.

- [027] A database system comprising the persistent storage system of claim 26.
- [028] An enterprise application comprising the persistent storage system of claim 26.
- [029] The system of claim 26, wherein the persistent storage structure comprises a database table.
- [030] The system of claim 26, wherein the persistent storage structure comprises a file directory.
- [031] The system of claim 26, wherein the access methods comprise methods for one of storing, retrieving, searching, and removing entity instance data in the persistent storage medium.
- [032] The system of claim 26, wherein the autonomous persistent storage system is an electronic catalog system.
- [033] The system of claim 26, wherein the autonomous persistent storage system further comprises an index creation module for automatically generating an index to persistent stored data.
- [034] The system of claim 26, wherein the autonomous persistent storage system further comprises a cache memory module for automatically storing an entity instance that is accessed from the persistent storage medium.
- [035] The system of claim 26, wherein the autonomous persistent storage system further comprises means for automatically populating a persistent storage structure with entity instance data that is input to the system.
- [036] An e-service that implements the system of claim 26 for providing a data management service based on a fee agreement or service level agreement.

[Fig.]




FIG. 1

[Fig. ]




FIG. 2

[Fig.]

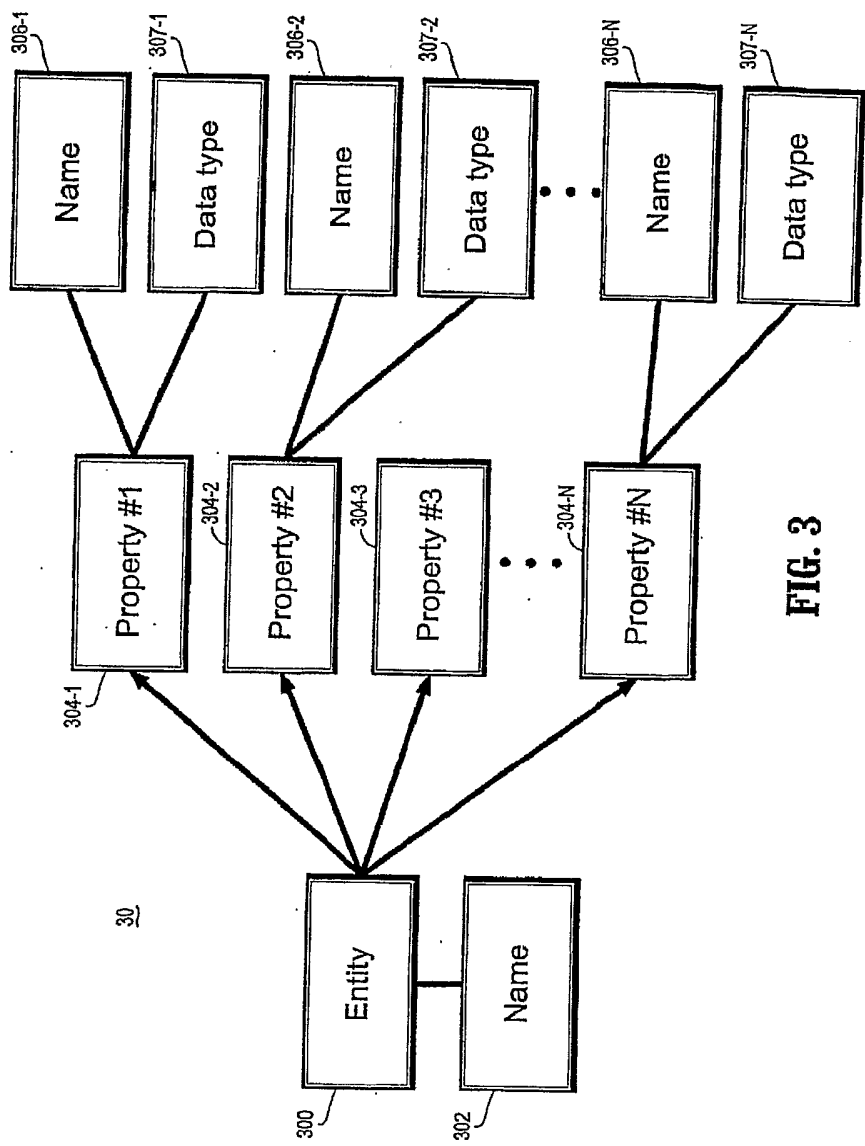


FIG. 3



[Fig. ]


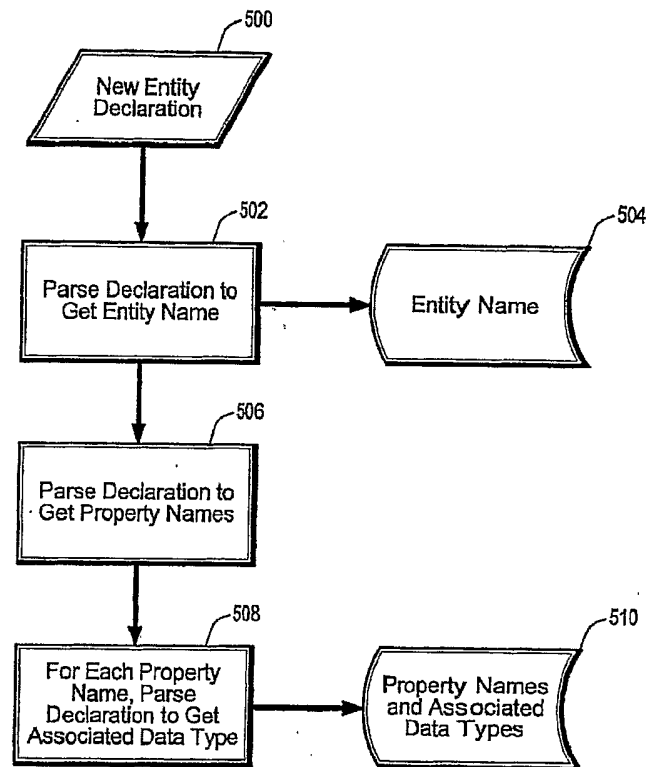


FIG. 4

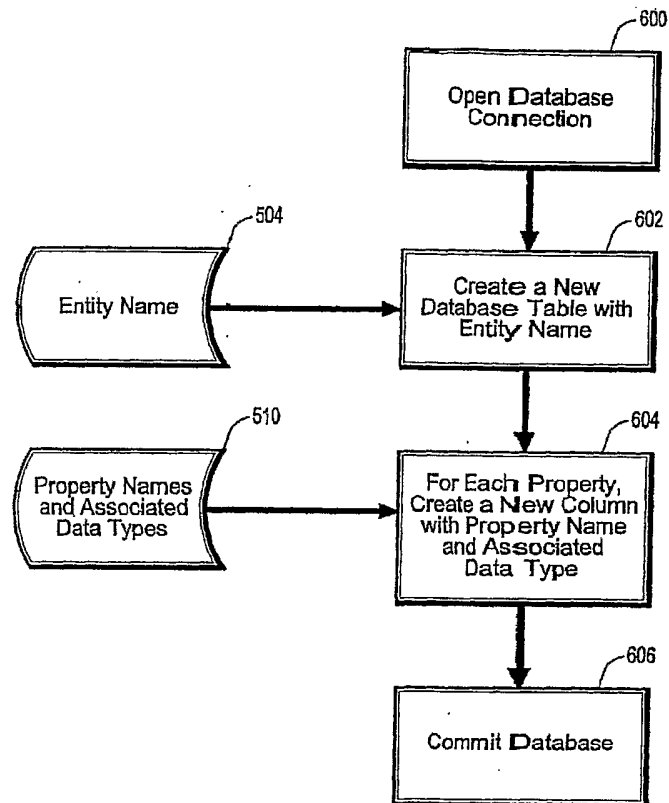
5/10

[Fig.]

**FIG. 5**

6/10

[Fig. ]

**FIG. 6**

7/10

[Fig. ]

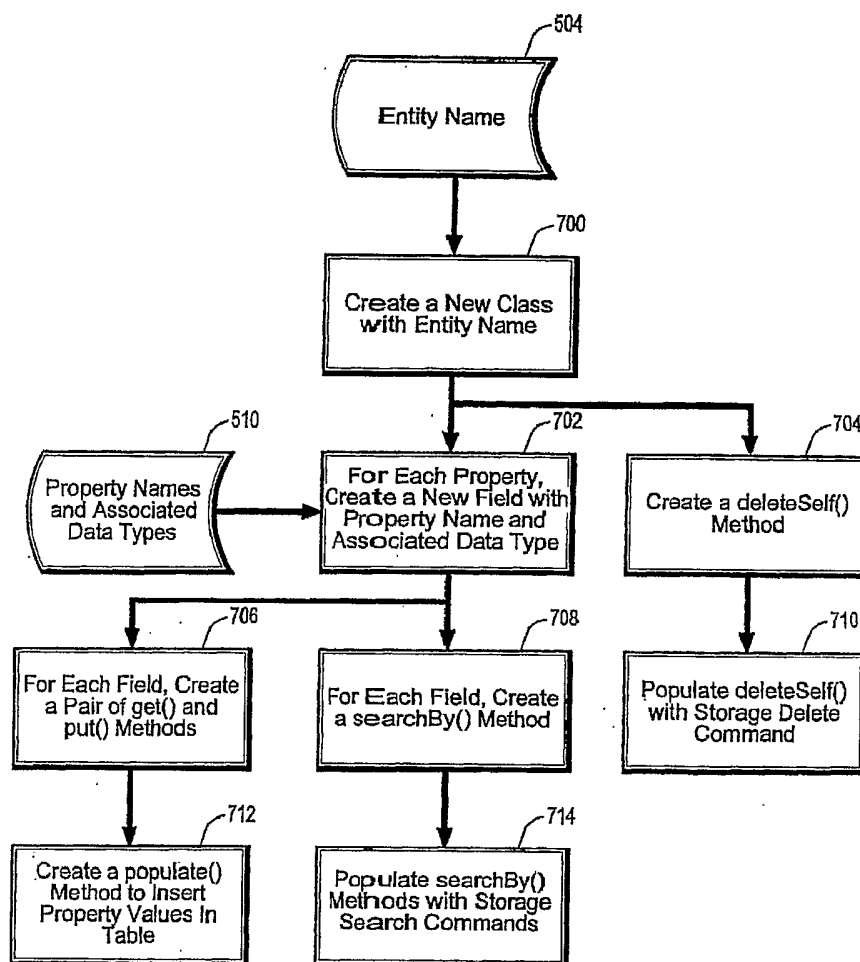


FIG. 7

8/10

[Fig. ]

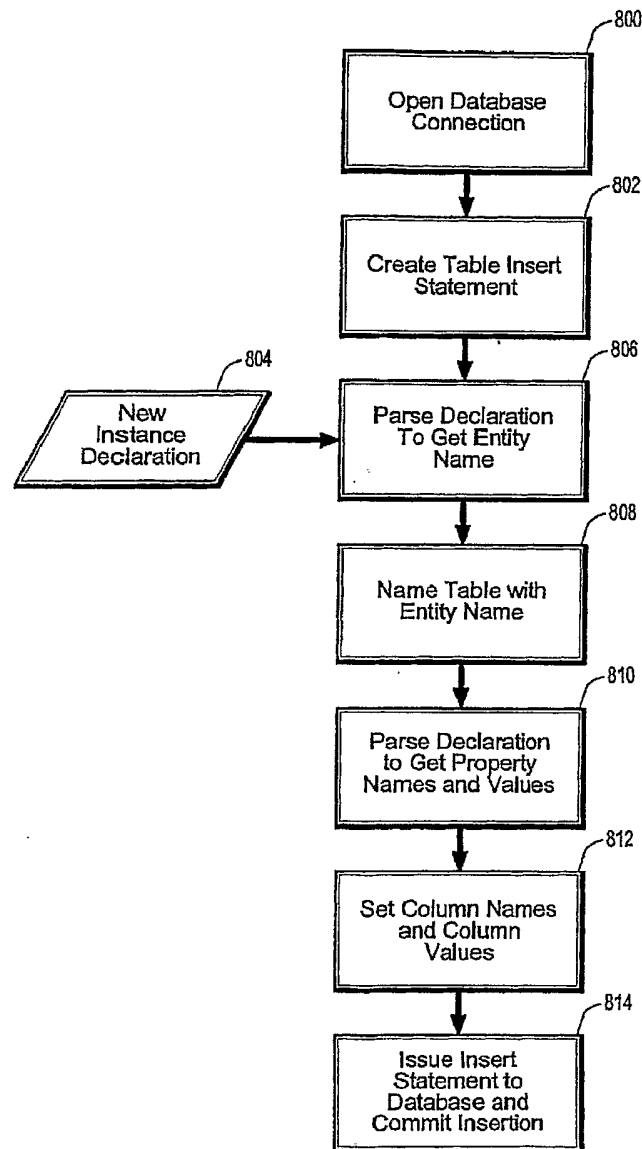
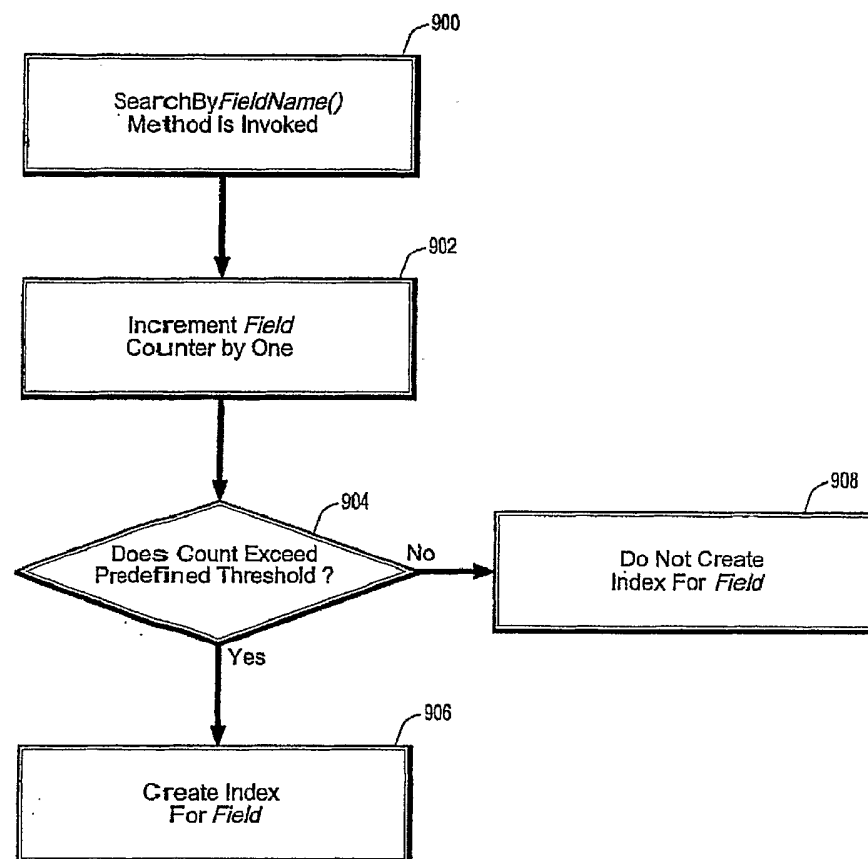


FIG. 8

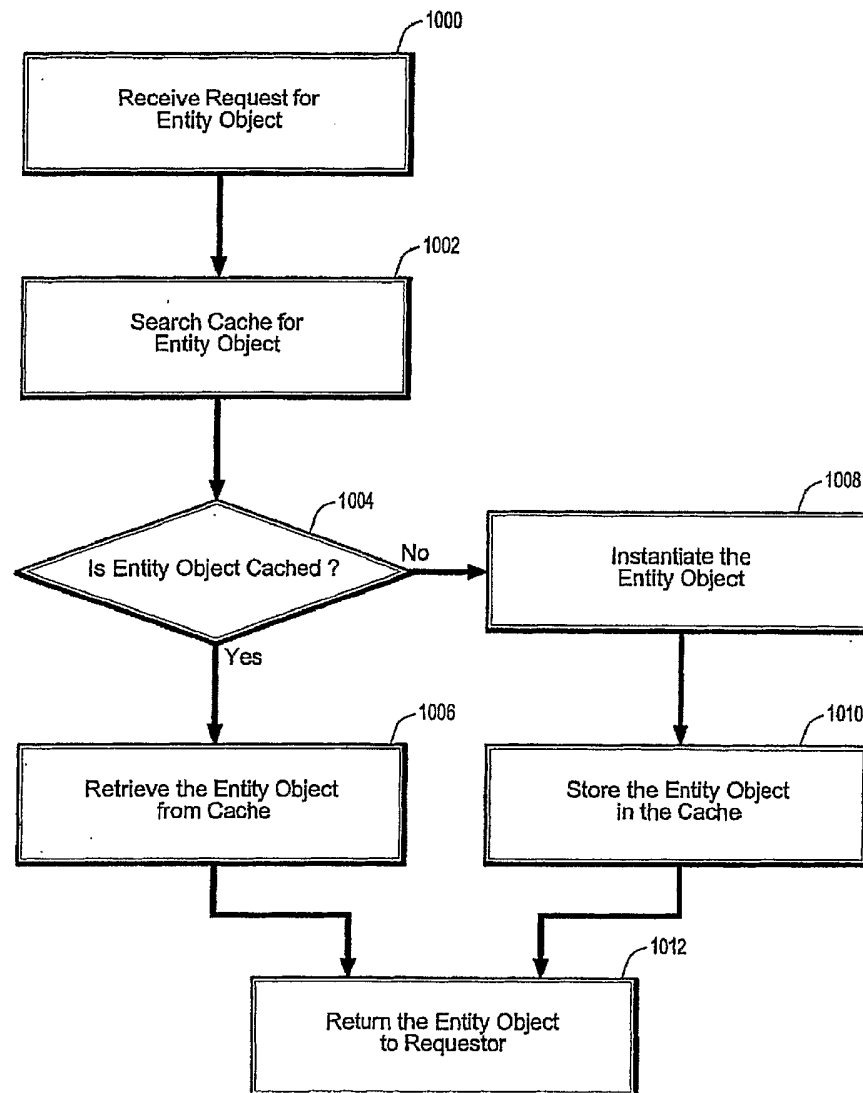
9/10

[Fig. ]

**FIG. 9**

10/10

[Fig. ]

**FIG. 10**

# INTERNATIONAL SEARCH REPORT

International Application No  
PCT/EP2004/052194

**A. CLASSIFICATION OF SUBJECT MATTER**  
IPC 7 G06F17/30

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

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)  
IPC 7 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

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 6 567 820 B1 (WALDO JAMES H ET AL) 20 May 2003 (2003-05-20) column 2, line 32 - line 63; figures 4,5,9A,9B column 4, line 10 - line 45 column 10, line 29 - column 11, line 15 column 5, line 65 - column 6, line 9 column 7, line 13 - column 9, line 2	1-36
A	US 5 870 742 A (HIGH JR ROBERT HOWARD ET AL) 9 February 1999 (1999-02-09) abstract ----- -/--	13,24,34

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

☒ Patent family members are listed in 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

5 January 2005

Date of mailing of the international search report

14/01/2005

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

Jaedicke, M



## INTERNATIONAL SEARCH REPORT

International Application No  
PCT/EP2004/052194

## C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	DANISH S: "Building Database-driven Electronic Catalogs" ACM SIGMOD RECORD, vol. 27, no. 4, December 1998 (1998-12), pages 15-20, XP002245952 cited in the application the whole document -----	1-36
A	BARTELS D ED - SU S Y W: "ODMG 93-the emerging object database standard" DATA ENGINEERING, 1996. PROCEEDINGS OF THE TWELFTH INTERNATIONAL CONFERENCE ON NEW ORLEANS, LA, USA 26 FEB.-1 MARCH 1996, LOS ALAMITOS, CA, USA, IEEE COMPUT. SOC, US, 26 February 1996 (1996-02-26), pages 674-676, XP010158969 ISBN: 0-8186-7240-4 the whole document -----	
A	US 5 295 256 A (BAPAT SUBODH) 15 March 1994 (1994-03-15) abstract -----	
A	GB 2 253 500 A (DATA GENERAL CORP) 9 September 1992 (1992-09-09) abstract -----	

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No  
PCT/EP2004/052194

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 6567820	B1	20-05-2003	US 6182083 B1	30-01-2001
			US 6032151 A	29-02-2000
			AU 2878399 A	15-09-1999
			CN 1298523 T	06-06-2001
			EP 1057123 A1	06-12-2000
			JP 2002505484 T	19-02-2002
			WO 9944157 A1	02-09-1999
			US 6708171 B1	16-03-2004
			US 6480863 B1	12-11-2002
			US 6578044 B1	10-06-2003
US 5870742	A	09-02-1999	KR 210208 B1	15-07-1999
US 5295256	A	15-03-1994	NONE	
GB 2253500	A	09-09-1992	NONE	