



(19) **United States**

(12) **Patent Application Publication**

(10) **Pub. No.: US 2002/0169767 A1**

(43) **Pub. Date: Nov. 14, 2002**

HARVEY

(54) **TABLE ARRANGEMENT FOR A DIRECTORY SERVICE SYSTEM AND FOR RELATED METHOD FACILITATING QUERIES FOR THE DIRECTORY**

(76) Inventor: **RICHARD HANS HARVEY,**
RINGWOOD (AU)

Correspondence Address:
SUGHRUE MION ZINN MACPEAK & SEAS PLLC
2100 PENNSYLVANIA AVENUE N W
WASHINGTON, DC 200373213

(*) Notice: This is a publication of a continued prosecution application (CPA) filed under 37 CFR 1.53(d).

(21) Appl. No.: **09/427,267**

(22) Filed: **Oct. 26, 1999**

Related U.S. Application Data

(62) Division of application No. 08/793,575, filed on May 22, 1997, now Pat. No. 6,052,681.

(30) **Foreign Application Priority Data**

Sep. 1, 1994 (AU)..... PM 7842
Nov. 21, 1994 (AU)..... PM 9586

Publication Classification

(51) **Int. Cl.⁷** **G06F 17/30**
(52) **U.S. Cl.** **707/4**

(57) **ABSTRACT**

A directory services system, such as one providing a X.500 or LDAP directory service, having a design and a method of operation which facilitates queries in and out of the directory, table organisation/layout and clustering. In the arrangement of the directory service, tables are organized corresponding to their function, thereby permitting an arrangement based on service modeling and functional organisation that permits clustering. A preferred layout design incorporates principal, conceptual, logical and/or physical designs.

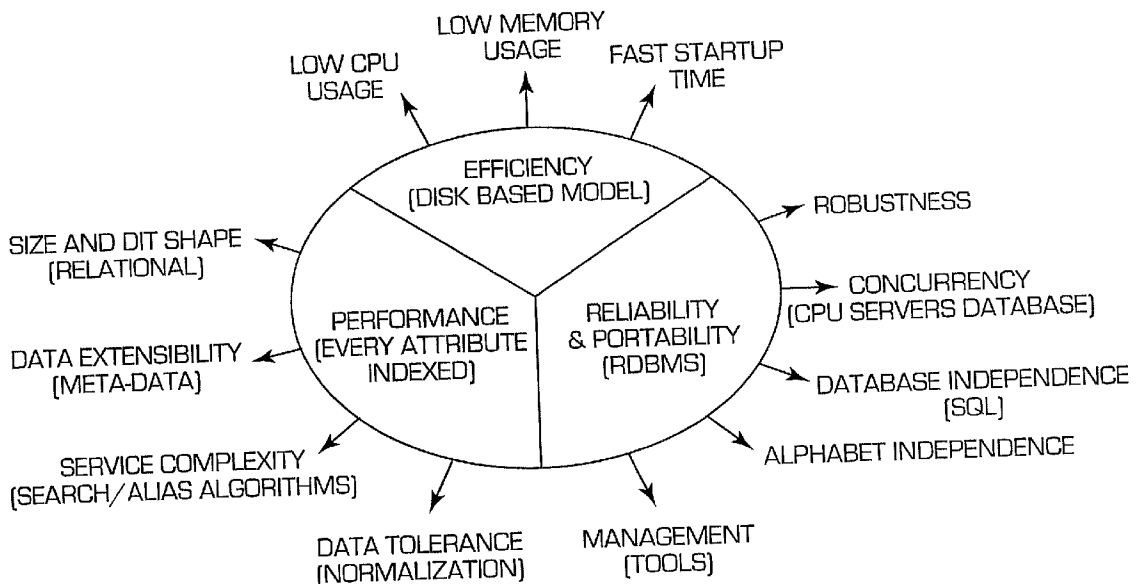


FIG. 1

OSI DIRECTORY PRODUCTS

3COM	1	0	1
ALCATEL TITN	1	1	1
BOLDON JAMES	1	0	1
BT	1	1	1
BRUNEL UNIVERSITY	1	0	1
CONTROL DATA SYSTEMS	1	1	1
CRAY RESEARCH	1	1	1
DATA CONNECTION	3	1	2
DIGITAL	2	1	2
GPT	1	1	1
HEWLETT PACKARD	1	1	1
IBM	5	5	5
ICL	5	4	4
ISODE CONSORTIUM	1	1	1
MARBEN PRODUIT	2	2	2
MOTOROLA COMP GROUP	1	1	0
NCR	1	1	1
NEXOR	3	1	2
OISWARE	1	1	1
OLIVETTI	1	1	1
RETIX	1	1	1
SIEMENS NIXDORF	2	2	2
SOFTWARE KENETICS	1	1	1
STRATUS COMPUTER	2	2	2
TANDEM COMPUTERS	2	2	2
UNISYS	5	5	5
WANG	1	1	1
WOLLONGONG GROUP	2	2	2

FIG. 2A

PRINCIPAL DESIGN

REPRESENTING X.500 IN A RDBMS

- DATA EXTENSIBILITY AND COMPLEXITY
- OBJECT ORIENTATED AND HIERARCHICAL

EMP#	NAME	AGE	SALARY
------	------	-----	--------

↓ RELATIONAL

TYPE	SYNTAX	VALUE
------	--------	-------

↓ X.500

PROPERTY

OBJECT NAME	PARENT NAME	TYPE	SYNTAX	VALUE
----------------	----------------	------	--------	-------

FUNCTIONAL DECOMPOSITION →

CONCEPTUAL DESIGN

IMPLEMENTING X.500 IN A RDBMS

- ATTRIBUTES AND VALUES
- HIERARCHY AND NAMES
- ALIASES
- DATA TOLERANCE

HIERARCHY

EID	PARENT	ALIAS	NAME
-----	--------	-------	------

↓

PARENT	PATH
--------	------

↓

ALIAS	A-EID
-------	-------

↓

NAMENORM	NAMERAW
----------	---------

OBJECT

EID	AID	VID	DISTING	VALUE
-----	-----	-----	---------	-------

↓

NAMENORM	NAMERAW
----------	---------

ATTRIBUTE

AID	TYPE	SYNTAX	OBJECTID
-----	------	--------	----------

SERVICE DECOMPOSITION →

FIG. 2B

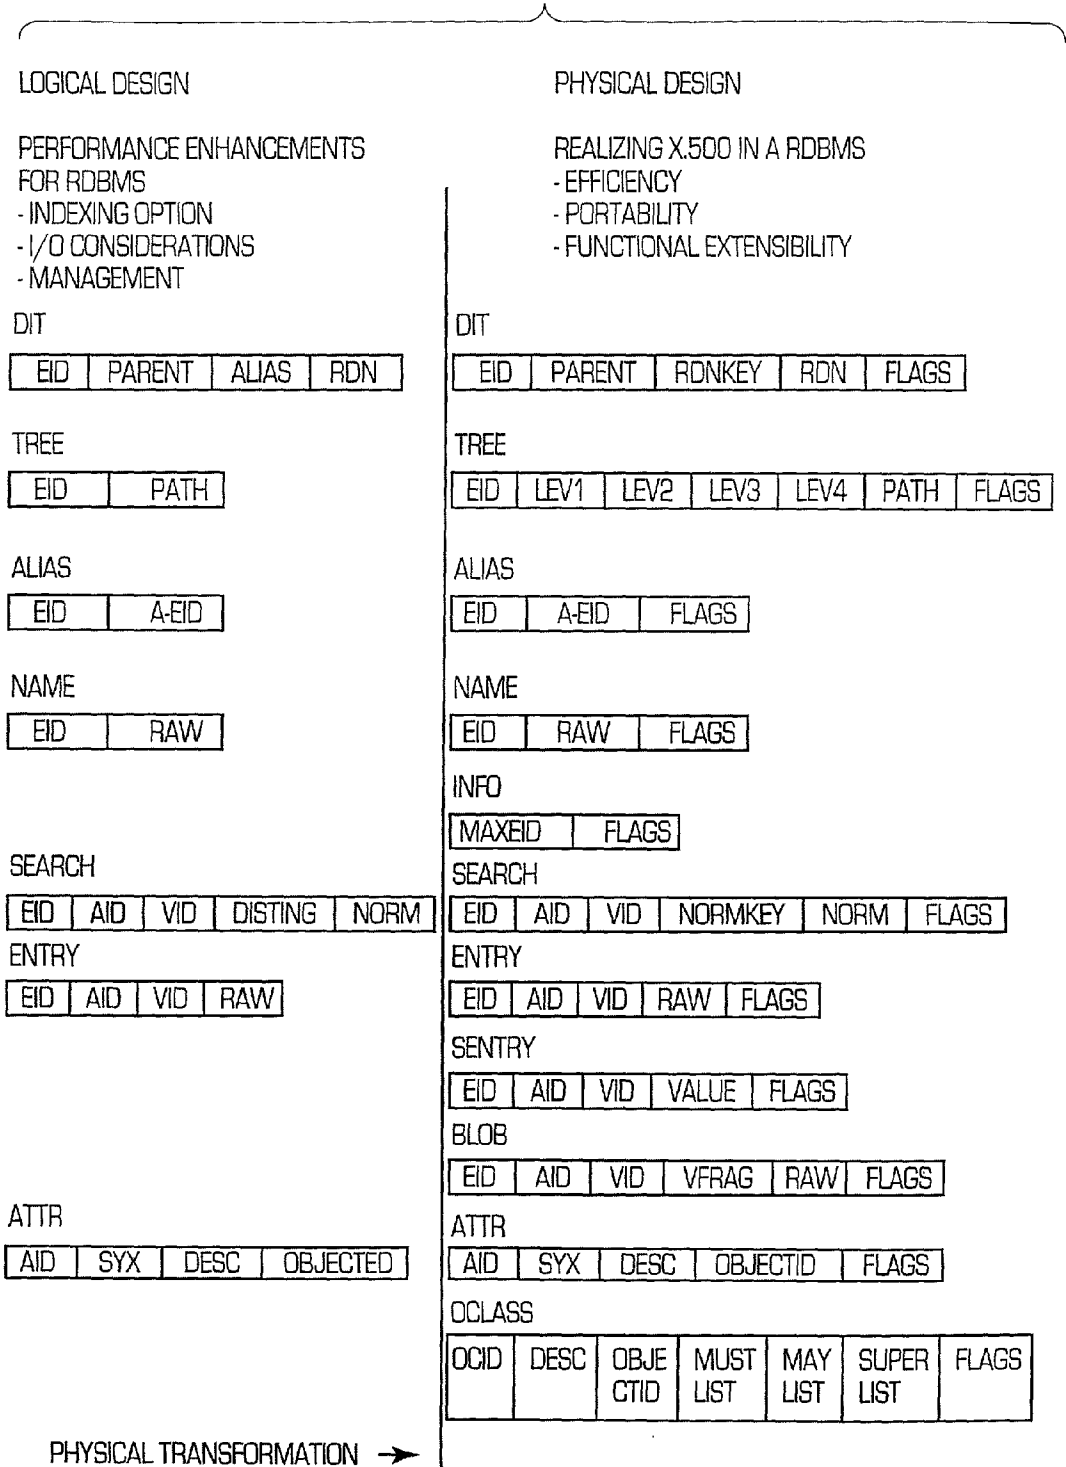


FIG. 3




FIG. 4

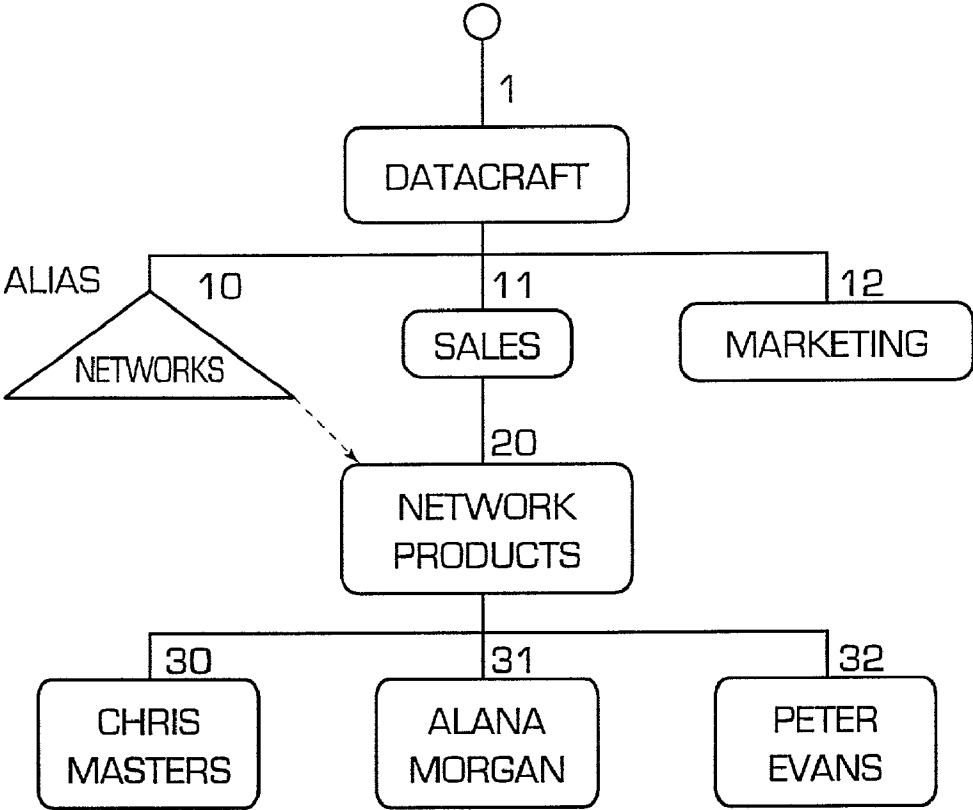


TABLE 3b/5a




TABLE ARRANGEMENT FOR A DIRECTORY SERVICE SYSTEM AND FOR RELATED METHOD FACILITATING QUERIES FOR THE DIRECTORY

[0001] This is a divisional of U.S. Ser. Number 08/793, 575, which is currently pending and which is incorporated herein by reference in its entirety.

BACKGROUND OF THE INVENTION

[0002] 1. Field of the Invention

[0003] The present invention relates to the field of directory services. In particular, the present invention is directed to application of X.500, LDAP and similar services to a relational database, a database design and use of the database to perform X.500 services.

[0004] One aspect of the invention relates to the design of a directory service(s) system, including table organisation and clustering, and method of operation which facilitates queries in and out of the directory.

[0005] Other aspects of the present disclosure are directed to an implementation using a RDBMS (Relational Database Management System) and also a table structure and methods of operation of a database application.

[0006] 2. Description of the Related Art

[0007] X.500 is the International Standard for Electronic Directories [CCITT89 or ITU93]. These standards define the services, protocols and information model of a very flexible and general purpose directory. X.500 is applicable to information systems where the data is fairly static (e.g. telephone directory) but may need to be distributed (e.g. across organisations or countries), extensible (e.g. store names, addresses, job titles, devices etc.), object oriented (i.e. to enforce rules on the data) and/or accessed remotely.

[0008] Relational Database Management System

[0009] (RDBMS) provide facilities for applications to store and manipulate data. Amongst the many features that they offer are data integrity, consistency, concurrency, indexing mechanisms, query optimisation, recovery, roll-back, security. They also provide many tools for performance tuning, import/export, backup, auditing and application development.

[0010] RDBMS are the preferred choice of most large scale managers of data. They are readily available and known to be reliable and contain many useful management tools. There is a large base of RDBMS installations and therefore a large amount of existing expertise and investment in people and procedures to run these systems, and so data managers are looking to use this when acquiring new systems. Most relational database products support the industry standard SQL (Structured Query Language).

[0011] There has also been a move towards Object Oriented systems which provide data extensibility and the ability to handle arbitrarily complex data items. In addition, many corporations and government departments have large numbers of database applications which are not interconnected. Data managers are looking for solutions which enable them to integrate their data, and to simplify the management of that data. X.500 and its associated standards provide a framework and a degree of functionality that enables this to be achieved. The fact that X.500 is an

international standard means that data connectivity can be achieved across corporations and between different countries.

[0012] The problem, therefore, is to address the need of data managers and implement X.500 with all the flexibility of object-oriented systems but using an SQL product so that it can achieve the scalability and performance inherent in relational systems coupled with the stability, robustness, portability and cost-effectiveness of current SQL products.

[0013] There have been a number of attempts of solving the above problem and over a considerable period of time. None of the attempts have resulted in a product which has proven to be commercially accepted by the market, and thus in the market place there is a long felt need yet to be addressed.

[0014] FIG. 1 shows an abstract from the "GOSIPNews" issue No. 4, dated April 1994 (Source: "Interoperability Products" distributed in Australia by the Centre for Open Systems) and which lists X.500 products currently available. None of these products use a SQL database as an underlying data store, and none of these products therefore address successfully the market need of implementing X.500 using an SQL RDBMS.

[0015] The Proceedings of IFIP WG6.6 International Symposium (ISBN: 0444 889 167) have published a paper presented by Francois Perruchond, Cuno Lanz, and Bernard Plattner and entitled "A Relational Data Base Design for an X.500 Directory System Agent". The Directory System disclosed, as with many prior art systems, is relatively slow in operation, particularly where the database is relatively extensive and is incomplete in its implementation of X.500, such as aliases, subsearch and entry information.

[0016] Another attempt is disclosed in the proceedings of IREE, ISBN 0909 394 253, proceedings Apr. 22-24, 1991 by C. M. R. Leung. In that disclosure, there is described a database scheme in which a single entry table holds detailed information about each directory object, and is also incomplete in its implementation of X.500.

[0017] This approach has been discredited by a number of text books and knowledge in the art, such as "Object-Oriented Modeling and Design" by J. Rumbaugh, et al, 1991, ISBN 0-13-630054-5, in which at paragraph 17.3.8 it is clearly stated that "putting all entities in the one table is not a good approach to relational database design".

[0018] As noted above, there have been a number of attempts made to address prior art problems, but none of the attempts have resulted in a product that has proven to be commercially accepted by the market. In particular, the prior art has been unable to address problems that are associated with queries in and out of the directory, table organisation and layout, and clustering.

SUMMARY OF INVENTION

[0019] An object of the present inventions is to address problems associated with the prior art relating to design of a directory service(s) system and method of operation which facilitates queries in and out of the directory, table organisation/layout and clustering, or at least one of the prior art problems.

[0020] The present invention provides a directory service arrangement in which tables are organised corresponding to their function. Preferably, a directory service is provided in which arbitrary data is processed using a fixed set of queries/services.

[0021] Furthermore, in the present invention each service is modeled, instead of each data type as seen in the prior art, and relationships are defined between each service, rather than between each data type as in the prior art. Implementation of service modeling using relational queries to satisfy directory services, such as X.500 services, enables the benefits of RDBMS to be optimally exploited.

[0022] Moreover, a directory service arrangement that is based on service modeling and functional organisation leads to clustering and a preferred layout design. In clustering, data of the same type and similar values are clustered in the same area of a storage medium, such as a disk. The number of disk pages that have to be accessed are then significantly reduced. The preferred, but not only, layout design that incorporates principal, conceptual, logical and/or physical designs, is subsequently described and illustrated in the accompanying tables and figures.

[0023] A further aspect of invention resides in providing a directory services database in which tables are arranged according to a conceptual design. The conceptual design includes at least attribute, object and hierarchy tables. The attribute table allows new attribute types to be defined by adding respective rows to this table. The object table defines the attributes within each object. The hierarchy table defines the relationship between objects.

[0024] A further aspect of invention is a method of arranging tables of a directory service, such as a X.500 or LDAP directory service, the method including:

[0025] decomposing tables around primary service relationships and deriving secondary services via joins.

[0026] Preferably, in applying service decomposition, at least one of the following consideration are made:

[0027] (a) Columns that have strong relationships are preferred to be kept together (to avoid unnecessary joins);

[0028] (b) If the number of significant rows in a given column is independent of the other related columns, then that given column is a candidate for a separate table.

[0029] (c) If a column is only used for locating information (input) or only used for returning results (output) then it is a candidate for its own table.

[0030] (d) If a column is used as a key for more than one service then it is preferred to be a primary key and therefore in its own table, since each table can have only one primary key.

[0031] (e) Keys are preferred to be unique or at least strong (non-repetitious).

[0032] Preferably, the hierarchy table can be decomposed into at least 4 tables, comprising a DIT (Directory Information Tree including information required for navigation), NAME (including information required for returning names

(raw RDN)), TREE (including information about an entry's path) and ALIAS (includes entries that are aliases), or their equivalent.

[0033] Preferably, the object table can be decomposed into at least 2 tables, namely SEARCH (used to resolve filters in a Search service and includes one row for each attribute value of each entry) and ENTRY (used to return values in Read and Search services and includes one row for each attribute value of each entry), or their equivalent.

[0034] Preferably, the invention includes a number of tables and designs having characteristics and relationships, as illustrated in **FIGS. 2A and 2B**.

[0035] A detailed description of the present invention can be found at least in the summary of invention, and in sections 1, 2, 4, particularly 4.1 and 4.6) and 6 of the description of the preferred embodiments, and the related Figures.

[0036] With regard to the remainder of the specification as a whole, in general, it seeks to disclose a number of other inventions related to the implementation of X.500 services in a RDBMS which supports SQL or any other relational language. X.500 services can be invoked via a number of protocols, such as X.500 and LDAP.

[0037] The scope of the present invention is outlined in this specification, including the claims.

[0038] In this document, at the time of filing, SQL is the most popular relational language and although it is only one form of relational language, the intent of the present invention is to have application to any other form of relational language, not just SQL.

[0039] These inventions can be related to the following headings:

1.	Principal Design
2.	Conceptual Design
3.	Conceptual Method(s)
4.	Logical Design
5.	Logical Method(s)
6.	Physical Design
7.	Example Implementation

[0040] The X.500 standard in no way dictates how the directory is to be implemented, only its capabilities and behaviour. One key to solving the implementation problem is the realisation that X.500 defines a fixed set of services (e.g. Add, Modify, Search etc.) that can operate on arbitrary data.

[0041] It has been discovered that problems associated with the prior art may be alleviated by a unique approach, by what may be described as inverting relational theory modeling from a data modeling approach to a service modeling approach. That is, from the problem of:

[0042] processing arbitrary queries on a fixed set of data to the present approach of processing arbitrary data using a fixed set of queries/services.

[0043] Each service is modelled (instead of each data type) and the relationships between each service defined (instead of the relationships between each data type).

[0044] Implementation of service modeling using relational queries to satisfy X.500 services enables benefits of RDBMS to be exploited.

[0045] The benefits of this approach are many. A summary is illustrated in FIG. 3. Some of the benefits include:

- [0046] relatively fast starting time.
- [0047] the ability to reduce memory requirements relative to memory resident systems.
- [0048] the ability to base X.500 on any SQL database and thereby protect the investment in products, expertise and procedures in managing existing systems.
- [0049] the ability to achieve performance relatively independent of size and relatively independent of the complexity of the data type. Every data type is treated generically. Every data type has an index on it. The result of indexing gives the ability to efficiently search the directory without caching large portions of directory into memory. Unlike the prior art where either only one index can be used to satisfy one given query or large portions of information is system intensively cached and searched in memory.
- [0050] the ability to support different languages (e.g. Spanish, Hebrew and Kanji) which may have various collating sequences. Single, double or other byte character sets may also be supported.
- [0051] using a disk based model to minimise I/O and efficiently retrieve I/O.
- [0052] the ability to service complex X.500 searches.
- [0053] the ability to create X.500 databases of far greater size than previously possible, without compromising performance or robustness. The databases can be small or large (250,000, 1 million or more entries).
- [0054] an optimal table design minimises wastage of disk space.
- [0055] the ability to leverage off hundreds of man years of relational database developments and use “industrial strength” databases with proven reliability, integrity, security and tools for developing high performance applications.

[0056] Based on this unique approach, the following disclosure will detail a number of inventions in an order with reference to FIGS. 2A and 2B, which illustrates schematically an overview of the present X.500 system. The table and column, names, order of columns and numeric values disclosed are given on an arbitrary basis in the overview. The number of columns disclosed represent a preferred operable requirement. Additional columns do not alter the use of the table as herein contemplated.

BRIEF DESCRIPTION OF THE DRAWINGS

[0057] FIG. 1 is an illustration of a table that lists X.500 products currently available, none of which use a SQL data base as an underlying data store.

[0058] FIG. 2A is an illustration schematically of an overview of the present invention, particularly the principal design and the corresponding conceptual design, as applied to the provision of a table structure for an X.500 system.

[0059] FIG. 2B is an illustration schematically of an overview of the present invention, particularly the logical design and the corresponding physical design, as applied to the provision of a table structure for an X.500 system.

[0060] FIG. 3 is an illustration of a pie chart that provides a summary representation of the benefits of implementing service modeling using relational queries to satisfy X.500 services.

[0061] FIG. 4 is an illustration of a hierarchy within a hypothetical organization, arranged as a tree, that is used to explain the services that may be provided according to the present invention.

[0062] FIG. 5 is an illustration of a hierarchy within a hypothetical organization, arranged as a tree, that has an alias referencing a different branch of the tree, according to the present invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0063] 1. Principal Design

[0064] The X.500 prior art attempts at implementation have been unable to overcome the relatively basic structural and operational differences between the X.500 requirements and functionality and SQL. The X.500 standard has a particular structure by nature, whereas SQL is designed to operate on relational structured tables.

[0065] For a typical relational database application, the nature of data is well known, i.e. tables will consist of a number of columns and each column contains data relating to a particular data type (see Table B1). The different data types that can be stored is limited to the columns of the table. The data types are also limited to the types supported by the database (e.g. string, numeric, money, date). The database may also store data of a form not understood by the database per se, but understood by the application e.g. binary data.

TABLE B1

Employee Table			
Name	Surname	Title	Phone
Chris	MASTERS	Sales Manager	03 727-9456
Alana	MORGAN	Sales Support	03 727-9455
...

[0066] If a new data type needs to be added (e.g. mobile) then a new column will have to be added to the table. This can cause problems if data table changes are not easy to implement. Also if the new data type is not well used (e.g. less than 1% of the organisation) then significant redundant data storage may result. See Table B2.

TABLE B2

<u>Employee Table</u>				
Name	Surname	Title	Phone	Mobile
Chris	MASTERS	Sales Manager	03 727-9456	018 042671
Alana	MORGAN	Sales Support	03 727-9455	
...

[0067] In essence, one invention in the application of X.500 resides in overcoming the extensibility by representing the X.500 attributes of the prior art:

[0068] empl # name age salary

[0069] as described above, as

[0070] type syntax value,

[0071] the latter representation being an extensible representation and is thus adapted to implementation with SQL. The latter representation is known as meta-data. The meta-data "value" may be binary.

[0072] A further development based on the above principal design is the adaption of the 'principal design' to X.500. This adaption has been realised by the provision of a 'property table', in which object name and parent name is added to the 'principal design'.

[0073] Further benefits accrue from the implementation disclosed above; including:

[0074] a. independence of complexity of filter—the implementation disclosed may utilise a query optimiser provided in SQL, and therefore there is no need to replicate a query optimiser in each proprietary database to which the present invention is applied,

[0075] b. independence of size—the implementation disclosed has the ability to be scaled,

[0076] c. independence of depth of tree—the implementation disclosed has hierarchy comparability,

[0077] d. performance—if index is put on the type column, then each and every type is indexed.

[0078] 2. Conceptual Design

[0079] The prior art has had difficulty in implementing X.500 as it has not been structured for extensibility, object oriented and hierarchy which are requirements of X.500.

[0080] This is addressed, in one form, by functionally decomposing the 'property table' and thus resulting in what is called the Conceptual Design.

[0081] The conceptual design resides in providing at least one of:

[0082] 1. Attribute table, where extensibility is addressed by allowing the definition of a new attribute type in this table by adding a row to the table;

[0083] 2. Object table, which defines the attributes within each object; and/or

[0084] 3. Hierarchy table, which defines the relationship between the objects.

[0085] In another invention, this problem is addressed by providing table structures in accordance with those disclosed in FIGS. 2A and 2B.

[0086] Yet further inventions reside in addressing problems of data tolerance by providing in the present X.500 system for the replacement of the 'value' column of the object table with value 'norm' and value 'raw' columns and/or replacing the RDN column in the hierarchy table with 'name norm' and 'name raw' columns.

[0087] Further, the difficulty in prior art of accommodating aliases is addressed in the present X.500 system by providing an 'alias' column in the hierarchy table. The 'alias' column is flagged to indicate that, that entry is an alias.

[0088] Further refinement may be provided by replacing the 'alias' column with alias and A-EID columns. The A-EID provides information about where the alias points.

[0089] Still further refinement may be provided by replacing the 'parent' column in the hierarchy table with 'parent' and 'path' columns.

[0090] The 'path' addresses the problem of implementing X.500 search, with aliases and subtrees. The 'path' has at least two unique properties: a) to determine the absolute position in the hierarchy; and b) it is used to determine if an entry is in a given subtree by its prefix.

[0091] 3. Conceptual Method

[0092] A number of unique methods of interrogating the conceptual design are disclosed in the detailed description following, including:

[0093] a) mapping the X.500 services into a sequence of SQL statements;

[0094] b) the search strategy is to apply the filter over the search area using the path or parent columns, and/or;

[0095] c) in dealing with aliases during navigation—where an alias points is cached in the A EID column;

[0096] d) in dealing with alias during search—find the unique set of base objects which define areas of the tree that need to be searched, and then apply b) above to each area of the tree.

[0097] A further invention is realised by using the attribute table for incoming data to find the AID from the X.500 object ID and outgoing data read from the database, vice versa. Furthermore, for any incoming distinguished name, it is navigated to its appropriate EID, then each search is performed as required by X.500.

[0098] Still furthermore, for a search, filter and subtree searches can be provided by a single pass resolution and using the path column. One invention is to utilise a 'path' field to simultaneously apply an arbitrary filter over an arbitrary subtree. The complications of aliases is handled by applying the above method to a uniquely resolved subtree.

[0099] Yet another unique method is to store the "path" of each entry as a string. Each path will then be prefixed by the path of its parent entry. This is useful for the filter in the search service.

[0100] 4. Logical Design

[0101] The logical design is based on a service decomposition of the conceptual design, though the realisation that X.500 service components are independent.

[0102] The advantages accruing from this include:

- [0103] 1. Reduces the number of indexes per table, as more tables are provided. It has been found that primary indexes are most efficient (speed, size) and secondary indexes may have large overheads (speed, size).
- [0104] 2. Enable data in tables to be clustered. Clustering occurs as a result of its primary key (storage structure) and thus data may be organised on disk around its key. E.g. for the 'search' table, surnames may be clustered together.
- [0105] 3. Management—smaller tables are easier to manage, e.g. faster to update indexes, collect statistics, audit, backup, etc.
- [0106] 4. Reduced I/O—speed improvements due to smaller rows, means more rows per page and thus operations perform less I/O's.

[0107] 5. Logical Methods

[0108] A number of unique methods of interrogating the logical design tables are disclosed in the detailed description following.

[0109] In addition, one method resides in caching the attribute table. Thus, (with the exception of initial loading) no SQL statements are issued to the database. In the present X.500 system, conversions are performed in memory. This provides a substantial speed advantage.

[0110] Further, validation is performed in memory which avoids database roll-back. Roll-backs are time and system consuming.

[0111] Still further, for the arbitrary filter, a dynamic SQL equivalent is built. This enables arbitrary complexity in X.500 searches.

[0112] Also for search results, the present system utilises set orientation queries of SQL to avoid 'row at a time' processing. Thus search results may be assembled in parallel in memory.

[0113] 6. Physical Design

[0114] New tables and new columns are introduced to overcome column width and key size restrictions and to achieve space optimisations.

[0115] The following text is a disclosure of embodiments of the inventions outlined:

[0116] 1. Principal Design

[0117] With reference to FIG. 2A, the principal design addresses the basic problem of representing the extensible, object oriented and hierarchical nature of X.500 in relational tables. In this section it will be disclosed (with examples) that the principal table design can be represented by a single table as shown in Table 1 below.

TABLE 1

X.500 Property Table				
object name	parent name	type	syntax	value

[0118] Throughout this and the following sections all column names and their positions in each table are arbitrary. The intent is to define what they contain and how they are used.

[0119] 1.1 Extensibility

[0120] For a typical relational database application, the nature of data is well known, i.e., tables will consist of a number of columns and each column contains data relating to a particular data type (see Table 1.1a). The table is self descriptive, i.e. the relations between data items is implied by being on the same row (this is the basis of relational theory).

TABLE 1.1a

Typical relational table			
name	surname	title	phone
Chris	MASTERS	Sales Manager	03 727-9456
Alana	MORGAN	Sales Support	03 727-9455
...

[0121] However, the above approach is not extensible because the number of different data types is limited to the number of columns of the table. If a new data type needs to be added (e.g. mobile phone number) then a new column will have to be added to the table (see Table 1.1b). Any application accessing this table will need to be updated to explicitly query it.

TABLE 1.1b

Relational table with an extra column				
name	surname	title	phone	mobile
Chris	MASTERS	Sales Manager	03 727-9456	018 042671
Alana	MORGAN	Sales Support	03 727-9455	
...

[0122] Other problems also exist in practice. If the new data type is not well used (e.g. less than 1% of the organisation has a mobile phone) then the table will be sparse (e.g. if a given person does not have a mobile then that row/column entry will be NULL). Also, the data types are limited to the types supported by the database (e.g. string, numeric, money, date, etc.).

[0123] The solution is to treat the data types as generic. The present invention adopts the method of representing arbitrary attributes (e.g. XOM [X/OPEN Object Management] API [Application Programming Interface]) as a type, syntax, value combination (see Table 1.1c)

TABLE 1.1c

Representing arbitrary attributes		
type	syntax	value
Name	String	Chris
Surname	String	MASTERS
Title	String	Sales Manager
Phone	Numeric	03 727-9456
Mobile	Numeric	018 042671

[0124] 1.2 Object Oriented

[0125] X.500 defines objects (e.g. people, organisations, etc.) which may contain an arbitrary number of “attributes”. Since many objects must appear in the table a mechanism is required to distinguish each object. An “object name” column is added to the table for this purpose (see Table 1.2a).

TABLE 1.2a

Representing objects with arbitrary values			
object name	type	syntax	value
Chris Masters	Name	String	Chris
Chris Masters	Surname	String	MASTERS
Chris Masters	Title	String	Sales Manager
Chris Masters	Phone	Numeric	03 727-9456
Chris Masters	Mobile	Numeric	018 042671
Alana Morgan	Name	String	Alana
Alana Morgan	Surname	String	MORGAN
Alana Morgan	Title	String	Sales Support
Alana Morgan	Phone	Numeric	03 727-9455

[0126] The above method allows any number of attributes to be assigned (related) to an entry. These attributes could be of arbitrary complexity (e.g. a multi-line postal address could be handled). As the number of columns is fixed new attributes can be added to any object without having to redefine the application. If a new attribute is added then an application that reads the entry will get back an extra row.

[0127] 1.3 Hierarchical

[0128] A method of representing hierarchical systems (e.g. parts explosion) is to use a parent/child combination (see Table 1.3a)

TABLE 1.3a

Parts explosion hierarchy	
parent	child
car	engine
car	fuel system
...	...
engine	carburettor
engine	pistons
...	...
carburettor	fuel valve
carburettor	air valve
...	...

[0129] X.500 defines its objects to be hierarchical. The relationships between objects follow a tree structure where each object has a parent object and each parent can have zero or more children. This relationship can be represented in a

general PROPERTY table by the addition of a “parent name” column, which is used to store the name of the parent object (see Table 1.3b).

FIG. 1.3b - X.500 Property Table

object name	parent name	type	syntax	value
Datacraft	root	Organisation	String	Datacraft
Datacraft	root	Address	Postal	PO Box 353
			Address	Croydon VIC
Chris Masters	Datacraft	Name	String	Chris
Chris Masters	Datacraft	Surname	String	MASTERS
Chris Masters	Datacraft	Title	String	Sales Manager
Chris Masters	Datacraft	Phone	Numeric	03 727-9456
Chris Masters	Datacraft	Mobile	Numeric	018 042671
Alana Morgan	Datacraft	Name	String	Alana
Alana Morgan	Datacraft	Surname	String	MORGAN
Alana Morgan	Datacraft	Title	String	Sales Support
Alana Morgan	Datacraft	Phone	Numeric	03 727-9455

[0130] Note that the root of the tree has no parent. Thus, if both Chris and Alana work for Datacraft and Datacraft is a child of the root then we can say that Chris and Alana are children of Datacraft and that Datacraft is the parent of Chris and Alana.

[0131] 2. Conceptual Design

[0132] In Section 1 it was shown that a single Property Table could represent the extensible, object oriented and hierarchical nature of X.500 (see Table 2a).

TABLE 2a

Property Table				
object name	parent name	type	syntax	value

[0133] With reference to FIG. 2A in this section it will be shown that full X.500 functionality can be represented by using three tables as shown below (see Table 2b and FIG. 2A).

TABLE 2b

Full Conceptual Design						
Hierarchy Table						
EID	Parent	Path	Alias	A_EID	NameNorm	NameRaw
Object Table						
EID	AID	VID	Disting	ValueNorm	ValueRaw	
Attribute Table						
AID		Type	Syntax		ObjectId	

[0134] The conceptual design addresses major problems with implementing full X.500 functionality in relational tables. As each major design issue is presented, examples are provided to illustrate the solution.

[0135] 2.1 Functional Decomposition

[0136] The Property Table (FIG. 2A) can be decomposed into separate tables that reflect the hierarchical, object oriented and extensible nature of X.500, preferably as follows;

- [0137] a Hierarchy Table which defines the structural relationship between objects.
- [0138] an Object Table which defines the attribute values within each object.
- [0139] an Attribute Table which defines the different attribute types.

[0140] These tables result from a process called functional decomposition.

[0141] To address the problem of correlating the relationships between tables, arbitrary numeric identifiers are introduced. The EID or “entry identifier” correlates each object with its hierarchy information. The AID or “attribute identifier” correlates each value in the object table with its attribute information.

[0142] The design is considered very efficient because the repeating groups in the Property table (type-syntax and object name-parent name) have been removed. Also, for SQL, the joining columns are simple integers.

TABLE 2.1		
Basic Conceptual Design		
<u>Hierarchy Table</u>		
BID	Parent	Name
10	0	Datacraft
30	10	Chris Masters
31	10	Alana Morgan
<u>Object Table</u>		
BID	AID	Value
10	10	Datacraft
10	16	PO Box 123 CROYDON
30	3	Chris
30	4	MASTERS
30	12	Sales Manager
30	20	03 727-9456
31	3	Alana
31	4	MORGAN
31	12	Sales Support
31	20	03 727-9455
<u>Attribute Table</u>		
AID	Type	Syntax
3	Name	string
4	Surname	string
10	Organisation	string
12	Title	string
16	Postal Address	address string
20	Phone	telephone string

[0143] 2.2 X.500 Attributes

[0144] X.500 attributes have a protocol identifier which is transferred when any data is communicated between end systems. These identifiers are internationally defined and are called OBJECT IDENTIFIERS (e.g. 2.5.4.4 means a sur-

name string). Thus an “ObjectId” column can be added to the Attribute table so that conversions between X.500 object identifiers and the internal attribute identifiers can be performed.

[0145] In addition, X.500 allows an attribute to have an arbitrary number of values (e.g. the mobile phone could be treated just as a second telephone number). Thus a “value identifier” or VID is introduced to identify values within an attribute in the Object Table.

TABLE 2.2			
Conceptual Design with X.500 attributes			
<u>Hierarchy Table</u>			
BID	Parent	Name	
10	0	Datacraft	
30	10	Chris Masters	
31	10	Alana Morgan	
<u>Object Table</u>			
BID	AID	VID	Value
10	10	1	Datacraft
10	16	1	PO Box 123 CROYDON
30	3	1	Chris
30	4	1	MASTERS
30	12	1	Sales Manager
30	20	1	03 727-9456
30	20	2	018 042671
31	3	1	Alana
31	4	1	MORGAN
31	12	1	Sales Support
31	20	1	03 727-9455
<u>Attribute Table</u>			
AID	Type	Syntax	ObjectId
3	Name	string	2.5.4.3
4	Surname	string	2.5.4.4
10	Organisation	string	2.5.4.10
12	Title	string	2.5.4.12
16	Postal Address	address string	2.5.4.16
20	Phone	telephone string	2.5.4.20

[0146] 2.3 X.500 Names

[0147] In X.500, each entry uses one or more of its attribute values (Distinguished Values) for naming the entry. A “Disting” column is added to the Object Table to flag the distinguished values.

[0148] The Distinguished Values combine to form a Relative Distinguished Name (RDN) which names the entry. The “Name” column in the Hierarchy table stores the RDN. This is an optimisation that negates the need for the RDN to be constructed from the distinguished values in the Object table.

[0149] An entry is uniquely named by a Distinguished Name (DN) which consists of all the RDN’s of the of its ancestors down from the root and the RDN of the object itself. An innovation is to add a “path” column to the Hierarchy table which defines the absolute position of the entry in the tree as a list of EID’s. The path has three important properties;

- [0150] 1) enables fast construction of DN’s, (the EID list defines all the RDN’s)

- [0151] 2) enables fast subtree searches (see Conceptual Methods),
- [0152] 3) it is independent of its DN (any of the RDN's in the DN can be renamed without affecting the path).

TABLE 2.3

Conceptual Design with X.500 attributes and names				
Hierarchy Table				
EID	Parent	Path	Name	
10	0	10	Datacraft	
30	10	10.30.	Chris, MASTERS	
31	10	10.31.	Alana, MORGAN	
Object Table				
BID	AID	VID	Disting	Value
10	10	1	1	Datacraft
10	16	1	0	PO Box 123 CROYDON
30	3	1	1	Chris
30	4	1	1	MASTERS
30	12	1	0	Sales Manager
30	20	1	0	03 727-9456
30	20	2	0	018 042671
31	3	1	1	Alana
31	4	1	1	MORGAN
31	12	1	0	SalesSupport
31	20	1	0	03 727-9455
Attribute Table				
AID	Type	Syntax		ObjectId
3	Name	string		2.5.4.3
4	Surname	string		2.5.4.4
10	Organisation	string		2.5.4.10
12	Title	string		2.5.4.12
16	Postal Address	address string		2.5.4.16
20	Phone	telephone string		2.5.4.20

[0153] 2.4 X.500 Aliases

[0154] X.500 also has the concept of ‘aliases’. An alias object effectively points to another entry and thus provides an alternate name for that entry. Thus an “alias” flag is added to the Hierarchy Table. When an alias is discovered during Navigation (i.e. the supplied DN contains an alias), then the alias value must be read from the Object Table. This alias DN must be resolved to where the alias points before Navigation of the original entry can continue.

[0155] An innovation is to use an “aliased EID” column or A_EID to store “where” the alias “points to”. This removes the need to repeatedly navigate through an alias.

TABLE 2.4

Conceptual Design with X.500 attributes, names and aliases					
Hierarchy Table					
EID	Parent	Path	Alias	A_EID	Name
10	0	10.	0	0	Datacraft
30	10	10.30.	0	0	Chris, MASTERS

TABLE 2.4-continued

Conceptual Design with X.500 attributes, names and aliases					
31	10	10.31.	0	0	Alana, MORGAN
35	10	10.35.	1	31	Support Engineer
<hr/>					
<div>Object Table</div> <hr/>					
EID	AID	VID	Disting	Value	
10	10	1	1	Datacraft	
10	16	1	0	PO Box 123 CROYDON	
30	3	1	1	Chris	
30	4	1	1	MASTERS	
30	12	1	0	Sales Manager	
30	20	1	0	03 727-9456	
30	20	2	0	018 042671	
31	3	1	1	Alana	
31	4	1	1	MORGAN	
31	12	1	0	Sales Support	
31	20	1	0	03 727-9455	
35	4	1	1	Support Engineer	
35	7	1	0	Datacraft/Alana,Morgan	
<hr/>					
<div>Attribute Table</div> <hr/>					
AID	Type	Syntax		ObjectId	
1	Alias Name	Distinguished Name		2.5.4.1	
3	Name	string		2.5.4.3	
4	Surname	string		2.5.4.4	
10	Organisation	string		2.5.4.10	
12	Title	string		2.5.4.12	
16	Postal Address	address string		2.5.4.16	
20	Phone	telephone string		2.5.4.20	

[0156] 2.5 X.500 Data Tolerance

[0157] Every X.500 attribute has a (internationally defined) syntax. X.500 attribute syntaxes define how each attribute should be treated. In all string syntaxes (e.g. Printable, Numeric etc.) superfluous spaces should be ignored. In some syntaxes the case is not important (e.g. Case Ignore String and Case Ignore List) and so the names “Chris Master”, “Chris MASTERS” and “ChRis MaSTeRS” are considered identical.

[0158] In order to do comparisons (e.g. search for a particular value), the syntax rules can be applied to create a normalised form (e.g. “CHRIS MASTERS”). If this normalised form is stored in the database, then any variations in input form are effectively removed, and exact matching can be used (which is necessary when using SQL).

[0159] Both the normalised data and “raw” data are stored in the database. The “raw” data is necessary so that users can retrieve the data in exactly the same format as it was originally input. As per the X.500 and LDAP standard, data received from a user, raw data, accords with ASN.1 (Abstract Syntax Notation No.1). Thus the “Name” column in the Hierarchy Table becomes the “NameRaw” and a “NameNorm” column is added. Similarly, the “Value” column in the Object Table becomes the “ValueRaw” and a “ValueNorm” column is added.

TABLE 2.5

Full Conceptual Design						
Hierarchy Table						
EID	Parent	Path	Alias	A_EID	NameNorm	NameRaw
10	0	10.	0	0	DATAACRAFT	Datacraft
30	10	10.30.	0	0	CHRIS,	Chris,
31	10	10.31.	0	0	MASTERS	MASTERS
					ALANA,	Alana,
35	10	10.35.	1	31	MORGAN	MORGAN
					SUPPORT	Support
					ENGINEER	Engineer
Object Table						
EID	AID	VID	Disting	ValueNorm	ValueRaw	
10	10	1	1	DATAACRAFT	Datacraft	
10	16	1	0	PO BOX 123	PO Box 123	
				CROYDON	CROYDON	
30	3	1	1	CHRIS	Chris	
30	4	1	1	MASTERS	MASTERS	
30	12	1	0	SALES	Sales Manager	
				MANAGER		
30	20	1	0	037279456	03 727-9456	
30	20	2	0	018321435	018 042671	
31	3	1	1	ALANA	Alana	
31	4	1	1	MORGAN	MORGAN	
31	12	1	0	SALES	Sales Support	
				SUPPORT		
31	20	1	0	037279455	03 727-9455	
35	4	1	1	SUPPORT	Support Engineer	
				ENGINEER		
35	7	1	0	DATAACRAFT /	Datacraft/Alana,	
				ALANA	Morgan	
				MORGAN		
Attribute Table						
AID	Type	Syntax		ObjectId		
1	Alias Name	Distinguished Name		2.5.4.1		
3	Name	Case Ignore String		2.5.4.3		
4	Surname	Case Ignore String		2.5.4.4		
10	Organisation	Case Ignore String		2.5.4.10		
12	Title	Case Ignore String		2.5.4.12		
16	Postal Address	Case Ignore List		2.5.4.16		
20	Phone	Telephone String		2.5.4.20		

[0160] 3. Conceptual Methods

[0161] This section introduces the basic X.500 services and shows how the conceptual table design, shown in Table

3a or FIG. 2A, is sufficient to implement X.500 services and their complexities.

TABLE 3a

Conceptual Table Design						
<u>Hierarchy Table</u>						
EID	Parent	Path	Alias	A_EID	NameNorm	NameRaw
<u>Object Table</u>						
EID	AID	VID	Disting	ValueNorm	ValueRaw	
<u>Attribute Table</u>						
AID		Type	Syntax		ObjectID	

[0162] The example hierarchy shown in Table 3b, as seen in FIG. 4, will be used to illustrate these services. Each name in the diagram represents an object entry in the database. The triangle represents an alias entry, and the dotted line represents the connection between the alias entry and the object that it points to. The numbers next to each entry are the entry EID's.

[0163] In the example, entry "1" has an RDN with a value of "Datacraft", entry "11" has an RDN with a value of "Sales", entry "20" has an RDN with a value of "Network Products" and entry "31" has an RDN with a value of "Alana Morgan". The DN of entry "31" is made up of a sequence of RDN's, namely, "Datacraft", "Sales", "Network Products", "Alana Morgan".

[0164] The alias entry "Datacraft/Networks" points to the entry "Datacraft", "Sales", "Network Products". When navigating to this entry the navigate process would find the alias entry, then find the DN of the object pointed to by the alias and then navigate from the root to the object entry returning an EID of "20" and a path of "1.11.20."

[0165] Listed below are sample tables which show how data is stored. The Hierarchy table (Table 3c) shows how the entries for the example hierarchy are stored. The Attribute table (Table 3e) shows attributes which are contained in the entry "Datacraft/Sales/Network Products/Chris Masters". The Object table (Table 3d) shows how the values of these attributes are stored.

TABLE 3

Sample Hierarchy Table						
EID	Parent	Path	Alias	A_EID	NameNorm	NameRaw
1	0	1.	0	0	DATAACRAFT	[Datacraft]
10	1	1.10.	1	20	NETWORKS	[Networks]
11	1	1.11.	0	0	SALES	[Sales]
12	1	1.12.	0	0	MARKETING	[Marketing]
20	11	1.11.20.	0	0	NETWORK	[Network
					PRODUCTS	Products]
30	20	1.11.20.30.	0	0	CHRIS MASTERS	[Chris
						Masters]

TABLE 3-continued

31	20	1.11.20.31.	0	0	ALANA MORGAN	[Alana Morgan]
32	20	1.11.20.32.	0	0	PETER EVANS	[Peter Evans]

Sample Object Table						
EID	AID	VID	Disting	ValueNorm	ValueRaw	
30	3	0	1	CHRIS	[Chris]	
30	4	0	1	MASTERS	[Masters]	
30	12	0	0	SALES MANAGER	[Sales Manager]	
30	20	0	0	03 727 9456	[(03) 727-9456]	
30	20	1	0	018 042 671	[(018) - 042 671]	

Sample Attribute Table				
AID	Type	Syntax		ObjectID
3	commonName	caseIgnoreString		2.5.4.3
4	surname	caseIgnoreString		2.5.4.4
12	title	caseIgnoreString		2.5.4.12
20	telephoneNumber	telephoneNumber		2.5.4.20

[0166] Distinguished Names

[0167] For the entry shown in the sample Object Table (Table 3d) two of the attributes, commonName and surname, are distinguished values (or naming values) which combine to form the RDN for the entry. This RDN is stored in the Hierarchy Table.

[0168] Multi-Valued Attributes

[0169] In X.500, it is permissible for an attribute to be multi-valued. The VID column is used to distinguish between values for an attribute. In the sample Object Table, the telephoneNumber attribute is multi-valued.

[0170] 3.1 Mapping Services to SQL

[0171] 3.1.1 Attribute Types and Values

[0172] Any data supplied by an X.500 service is supplied as a list of ObjectID's and their associated values. These must be converted into AID's (using the Attribute table) and normalised values (using the Object table) for use by the X.500 application. The database returns data as AID's and Raw Values, which must then be converted into ObjectID's and their associated values in the X.500 result.

[0173] 3.1.2 Navigation

[0174] Each X.500 service supplies a Distinguished Name which is converted into an EID for use by the X.500 application. When the application processes a service it returns one or more EID's. These EID's can then be translated back into Distinguished Names in the X.500 result.

[0175] All X.500 services rely on navigating the directory tree. To navigate to a particular entry, the following procedure is performed:

[0176] Given the DN for the entry, locate the entry in the hierarchy table which has an RDN equal to the first RDN in the DN.

[0177] Store the EID.

[0178] Recursively, locate the entry which has an RDN equal to the next RDN in the DN and a parent equal to the stored EID.

EXAMPLE

[0179] Navigate to the entry "Datacraft/Sales/Network Products/Peter Evans". This will result in a number of select statements, with each returned EID being used as the value of the PARENT in the next statement.

[0180] select EID from HIERARCHY

[0181] where PARENT=0 and RDN="DATA-CRAFT"

[0182] select EID from HIERARCHY

[0183] where PARENT=1 and RDN="SALES"

[0184] select EID from HIERARCHY

[0185] where PARENT=11 and RDN="NETWORK PRODUCTS"

[0186] select EID from HIERARCHY

[0187] where PARENT=20 and RDN="PETER EVANS"

[0188] 3.1.3 Read

[0189] Selected attributes to be read can be supplied. Only the values of these attributes (if they are present in the entry) will be returned.

[0190] 'Types only' can be selected as a read option, in which case no values will be returned. All types present in the entry, or those selected, will be returned.

[0191] Navigate to the entry to be read. Store the EID. In the Object Table, read the values of all rows which match the stored EID.

EXAMPLE

[0192] Read the entry "Datacraft/HQ/Network Products" and return all types and values.

[0193] Navigate to the entry (as in 3.1.2) and then;

[0194] select AID, VALUERAW from OBJECT

[0195] where EID=20

[0196] 3.1.4 Compare

[0197] Compare returns a 'matched' or 'not matched' result. A raw value is input but the compare is performed using the normalised value.

[0198] Navigate to the required entry. Store the EID. In the Object Table, test for a matching value in all rows which match the stored EID and the specified AID.

EXAMPLE

[0199] Compare the telephone number "03 727 9256" with the entry "Datacraft/Sales/Network Products/Chris Masters".

[0200] Navigate to the entry and then;

[0201] select VALUERAW from OBJECT

[0202] where EID=30

[0203] and AID=20

[0204] and VALUENORM="03 727 9456"

[0205] If a value is selected then return "matched" else return "not matched".

[0206] 3.1.5 List

[0207] Navigate to the required entry. Store the EID. In the Hierarchy Table, return the RDN's for all rows with a parent matching the stored EID.

EXAMPLE

[0208] List from the entry "Datacraft/Sales".

[0209] Navigate to the entry and then;

[0210] select NAMERAW from HIERARCHY

[0211] where PARENT=11

[0212] 3.1.6 Add Entry

[0213] Navigate to the required parent entry. Store the EID of the parent. Add a new EID to the Hierarchy table and add rows to the Object table for each value in the new entry.

EXAMPLE

[0214] Add a new entry under the entry "Datacraft/Sales/Network Products".

[0215] Navigate to the entry and then;

[0216] insert into OBJECT

[0217] (EID, AID, VID, DISTING, VALUENORM, VALUERAW)

[0218] values (33, 3, 1, 1, EDWIN MAHER, Edwin Maher) and

[0219] insert into HIERARCHY

[0220] (EID, PARENT, PATH, ALIAS, A-EID, NAMENORM, NAMERAW)

[0221] values (33, 20, 1.11.20.33.,0, 0, EDWIN MAHER, Edwin Maher)

[0222] 3.1.7 Remove Entry

[0223] Navigate to the required entry. Check that the entry is a leaf on the tree, (i.e. check that it has no subordinate

entries on the tree). Store the EID. Remove the entry from the Hierarchy table. In the Object Table, remove all rows which match the stored EID.

EXAMPLE

[0224] Remove an entry (with EID=33) under the entry "Datacraft/Sales/Network Products".

[0225] Navigate to the entry and then;

[0226] delete from OBJECT

[0227] where EID=33 and

[0228] delete from HIERARCHY

[0229] where EID=33

[0230] 3.1.8 Modify Entry

[0231] Navigate to the required entry. Store the EID. In the Object Table, Add, Remove or Modify rows matching the stored EID.

EXAMPLE

[0232] Modify the entry "Datacraft/Sales/Network Products/Alana Morgan".

[0233] Add value=title="Branch Manager".

[0234] Navigate to the entry and then;

[0235] select EID, AID, VID, VALUENORM from OBJECT

[0236] where EID=31

[0237] Test the returned rows for an attribute of title. If none exist, the attribute can be added, otherwise the attribute must be checked to see if it can be multi-valued and whether it already exists.

[0238] Insert into OBJECT

[0239] (EID, AID, VID, DISTING, VALUENORM, VALUERAW)

[0240] values (31,12,1,0, BRANCH MANAGER, Branch Manager).

[0241] 3.1.9 Modify RDN

[0242] Navigate to the required entry. Check that the new name (RDN) does not exist in the current level of the subtree (i.e. that the new DN is distinct). Store the EID. Modify the entry in the Hierarchy and Object tables.

EXAMPLE

[0243] Modify the RDN of the entry "Datacraft/Sales/Network Products/Chris Masters" to "Christine Masters".

[0244] Navigate to the entry and then;

[0245] select EID from HIERARCHY

[0246] where PARENT=20

[0247] and VALUENORM="CHRISTINE MASTERS"

[0248] If no entries are returned then the new RDN may be inserted. First set the old RDN to be a non-distinguished value.

- [0249] update OBJECT
- [0250] set DISTING=0
- [0251] where EID=30 and VALUENORM="CHRIS"
and
- [0252] update HIERARCHY
- [0253] set NAMENORM="CHRISTINE MAS-
TERS" and
- [0254] set NAMERAW="Christine Masters"
- [0255] where EID=30 and
- [0256] insert into OBJECT
- [0257] (EID, AID, VID, DISTING, VALUENORM,
VALUERAW)
- [0258] values (30, 3, 1, 1, "CHRISTINE", "Chris-
tine")

[0259] 3.2 Search Strategy

[0260] The most powerful and useful X.500 service is the search service. The search service allows an arbitrary complex filter to be applied over a portion of the Directory Information Tree (the search area).

[0261] A filter is a combination of one or more filter items connected by the operators AND, OR and NOT. For example; surname="MASTERS" AND title="SALES MANAGER"

[0262] The Search area is the part of the tree that is covered by the scope of the search (base-object-only, one-level or whole-subtree).

[0263] One technique for resolving searches is to apply the filter and then to see if any matching entries are in the search area. In this case a filter is applied to the entire tree and EID's for all rows matching the filter are returned. Then, for each EID found, step search up through the hierarchy to see if the entry is a subordinate of the base object (i.e. the entry has a parent/grandparent/ . . . that is the base object). If the number of matches is large and the subtree small this is very inefficient. This technique doesn't cope with aliases as an alias is not a parent of the object that it points to and many aliases may point to a single object.

[0264] A second strategy is to obtain a list of all EID's in the search area and then apply the filter to these EID's. If an alias is resolved that points outside of the original search area then the subtree pointed to by the alias is expanded and the EID's in that subtree are added to the list. The filter is then applied to the set of expanded EID's. This is very poor if the search area is large.

[0265] An innovation is to simultaneously apply the filter over the search area (instead of sequentially as in the two methods described above). This is called single pass resolution. This method is considered to provide considerable performance improvement over the above methods because the rows that are retrieved are those that satisfy both the filter and scope requirements of the search.

[0266] When performing a one level search the filter is applied to all entries that have a parent equal to the EID of the base object (for example; search where parent=20 will apply the filter to entries 30,31 and 32).

[0267] When performing a subtree search the path is used to expand the search area. The "path" of each entry is a string of numbers (e.g. "1.10.50.222." which indicates that entry

222 has a parent of 50, a grandparent of 10 and a great grandparent of 1). The path has the unique property that the path of an entry is a prefix of the path of all entries that are subordinate to the entry. That is the path of an entry forms the prefix of the paths of all entries in the subtree below the entry. Therefore when performing a subtree search we obtain the base object of the subtree and then apply the filter to all entries that have a path which is prefixed by the path of the base object (for example; to search for all entries under "Sales" we perform a search where PATH LIKE 1.11.%).

[0268] Base Object Search:

[0269] Navigate to the base object. Store the EID. In the Object Table, read nominated values from rows which match the stored EID where a filter criteria is satisfied, eg, telephone prefix="727".

EXAMPLE

[0270] Search from the base object "Datacraft/Sales/Network Products" for an entry with surname="MORGAN", using a "base-object-only" search. Navigate to the base object and then;

[0271] select AID, VALUERAW from OBJECT

[0272] where EID=20 and AID=4

[0273] and NAMENORM="MORGAN"

[0274] One Level Search:

[0275] Navigate to the base object. Store the EID. Return the list of EID's which have a parent EID matching the stored EID (in Hierarchy table) and have values which satisfy the filter criteria (OBJECT table). In the Object Table, read nominated values for the returned EID's.

EXAMPLE

[0276] Search from the base object "Datacraft/Sales/Network Products" for an entry with surname="MORGAN", using a "one-level-only" search. Navigate to the base object and then;

[0277] select H.EID from HIERARCHY H,
OBJECT O

[0278] where PARENT=20 and AID=4 and
NAMENORM="MORGAN"

[0279] and H.EID=O.EID

[0280] then place the EID's returned into an
EIDLIST and

[0281] select AID, VALUERAW from OBJECT

[0282] where EID in [EIDLIST]

[0283] Subtree Search:

[0284] Navigate to the base object. Store the EID. Return the list of all EID's with a path like that of the base object (Hierarchy table) and have values which satisfy the filter criteria (OBJECT table). In the Object Table, read nominated values for the returned EID's.

EXAMPLE

[0285] Search from the base object "Datacraft/Sales/Network Products" for an entry with surname="MORGAN", using a "whole-subtree" search. Navigate to the base object and then;

[0286] select H.EID from HIERARCHY H,
OBJECT O
[0287] where PATH like "1.11.20.%" and AID 4
[0288] and NAMENORM="MORGAN"
[0289] and H.EID=O.EID
[0290] then place the EID's returned into an
EIDLIST and
[0291] select AID, VALUERAW from OBJECT
[0292] where EID in [EIDLIST]

[0293] 3.3 Aliases and Navigate

[0294] Aliases are resolved during navigation if the "don't-dereference-alias" flag is not set and the service is not an update service (add, delete, modify, modifyRDN).

[0295] When an alias is discovered during navigation the alias must be resolved. That is, the object that the alias points to must be obtained. First we check the A_EID column of the Hierarchy table. If the A_EID is 0 then the object that the alias points to must be obtained from the Object table and this object must then be navigated to and the resultant EID stored in the A_EID column. If this is done successfully then the remainder of the path can be navigated. By storing the EID of the aliased object in the A_EID column of the Hierarchy table it is possible to avoid navigating to aliased objects. This can save time, especially if the aliased object is at a low level of the hierarchy.

[0296] 3.4 Aliases and Search

[0297] Aliases are dereferenced during a search if the "search-aliases" flag in the search argument is set. The performance of the search service while dereferencing aliases becomes a two step process. Firstly, define the search area and then apply the filter to the entries within the search area. Aliases dereferenced as part of the search service can expand the search area to which the filter is applied. They also restrict the search area in that any dereferenced aliases are excluded from the search area.

[0298] Aliases and OneLevel Search

[0299] If aliases are being dereferenced as part of a one level search and an alias entry is found then the alias must be resolved (using the Object table or the A_EID). The aliased object is then added to the search area to which the filter is applied. In a oneLevel search where aliases are found the search area will consist of non-alias entries directly subordinate to the base object and all dereferenced aliases.

[0300] Aliases and Subtree Search

[0301] If aliases are being dereferenced as part of a whole subtree search and an alias entry is found then the alias must be resolved (using the Object table or the A_EID) and this EID must then be treated as another base object, unless it is part of an already processed sub tree.

[0302] When dereferencing aliases during a search the "Path" column can be used to find alias entries within a subtree join. If an alias entry is found that points outside of the current subtree then the subtree pointed to by the alias can also be searched for aliases. One property of the hierarchical tree structure is that each subtree is uniquely represented by a unique base object (i.e. subtrees do not overlap). When performing a subtree search we build up a list of base objects which define unique subtrees. If no aliases are found then the list will contain only one base

object. If an alias is found that points outside of the subtree being processed then we add the aliased object to the list of base objects (unless one or more of the base objects are subordinate to the aliased object in which case the subordinate base object(s) are replaced by the aliased object). The search area will therefore consist of non-alias entries that have a path prefixed by the path of one of the base objects.

[0303] 4. Logical Design

[0304] Whilst the Conceptual Design (see Table 4a) is sufficient to implement the X.500 functionality, further performance improvements can be made.

TABLE 4a

Conceptual Design						
Hierarchy Table						
EID	Parent	Path	Alias	A_EID	NameNorm	NameRaw
Object Table						
EID	AID	VID	Disting		ValueNorm	ValueRaw
Attribute Table						
	AID	Type		Syntax		ObjectId

[0305] Performance improvements in conventional relational design can be achieved because assumptions can be made about the data—the data is essentially fixed at the time an application is designed. In X.500, none of the data types are known. However performance improvements can still be made because assumptions can be made about the services—these are known at the time the X.500 application is designed.

[0306] With reference to FIG. 2B, one innovative approach is to recognise that each table can be organised around the major service relationships (instead of around the major data relationships in conventional relational design). It shall be shown that the above tables can be decomposed into a number of smaller and more efficient tables as shown below.

TABLE 4b

Logical Design				
DIT	PARENT	ALIAS		RDN
EID NAME				
EID TREE		RAW		
EID ALIAS	PATH			
EID SEARCH			A_EID	
EID ENTRY	AID	VID	DISTING	NORM
EID ATTR	AID	VID		RAW
AID	SYNTAX	DESC	OBJECTID	

[0307] 4.1 Service Decomposition

[0308] The practical reality for most RDBMS's is that big tables with many columns do not perform as well as smaller tables with fewer columns. The major reasons are to do with indexing options, I/O performance and table management

(see Sections 4.5 and 4.6). This is why prior art relational design techniques aim to focus primary information into separate tables and derive secondary information via table joins (i.e. normalisation and fragmentation techniques).

[0309] One innovation in achieving X.500 performance is to decompose the tables around primary service relationships and derive secondary services via joins. This process is called service decomposition. The following considerations are made:

[0310] (1) Columns that have strong relationships are preferred to be kept together (to avoid unnecessary joins);

[0311] (2) If the number of significant rows in a given column is independent of the other related columns, then that given column is a candidate for a separate table.

[0312] (3) If a column is only used for locating information (input) or only used for returning results (output) then it is a candidate for its own table.

[0313] (4) If a column is used as a key for more than one service then it is preferred to be a primary key and therefore in its own table (each table can have only one primary key).

[0314] (5) Keys are preferred to be unique or at least strong (non-repetitious).

[0315] A first level analysis of column usage is shown in Table 4.1.

TABLE 4.1

Basic column usage												
X.500 Service	Table	EID	AID	VID	Value Norm	Value Raw	Parent	Alias	Name Norm	Name Raw	Path	
Navigate	H	R					S	R	S		R	
	O	S	(S)/R	R		R				R	R	
Compare	O	S	S		S							
	H						S	R		R		
List	O	S/R	S		(S)		(S)				(S)	
Search - filter												
Search - result		S/R	(S)/R	R		R				R	R	
Add	H/O	S										
Remove	H/O	S										
Modify	O	S	S	S	S							
Modify RDN	H/O	S	S	S					S			

[0316] Key to symbols in the above table:

[0317] H—Hierarchy table

[0318] O—Object table

[0319] S—Supplied value (used in the SQL for Searching the table)

[0320] R—Returned value (value retrieved from the tables)

[0321] ()—item may or may not be present depending on the options of the service.

[0322] From the above information and further analysis, the Conceptual Design tables can be decomposed into a number of smaller tables as described in the following sections.

[0323] 4.2 Hierarchy Table Decomposition

[0324] The Hierarchy table contains the following columns:

TABLE 4.2a

Hierarchy Table						
EID	Parent	Path	Alias	A_EID	NameNorm	NameRaw

[0325] The Hierarchy Table contains information about objects and their parents, their names, their absolute positions in the hierarchy and if they are aliases. This table can therefore be split into four tables: DIT, NAME, TREE and ALIAS.

[0326] The parent information is used for finding a given child or acting on entries that have a given parent. Finding a given child (e.g. Parent=0, NameNorm="DATAACRAFT") is the basis for Navigation and update checking (checking for the existence of an object before an Add or ModifyRdn). Acting on entries that have a given parent is used during List or OneLevel Search. Thus the DIT (Directory Information Tree) table has information required for Navigation, but allows its PARENT column to be used by other services.

TABLE 4.2b

DIT Table			
EID	PARENT	ALIAS	RDN

[0327] An object is differentiated from its siblings via its Relative Distinguished Name (RDN). RDN's are returned for a List (in conjunction with a given Parent) or as part of a full Distinguished Name (Read, Search). Thus the NAME table has information required for returning names (the raw RDN).

TABLE 4.2c

NAME Table	
EID	RAW

[0328] An object’s absolute position in the hierarchy is necessary for building DN’s (from which the raw RDN’s are retrieved) and for expanding subtrees during Search. Thus the TREE table has information about an entry’s Path (the sequence of EID’s down from the root).

TABLE 4.2d

TREE Table	
EID	PATH

[0329] Alias information is cached so that every time an alias is encountered during Navigate it does not have to be repeatedly resolved. Thus the ALIAS table only contains entries that are aliases. It is also used during OneLevel Search (in conjunction with the DIT Parent column) and Subtree Search (in conjunction with the Path column) to determine if there are any aliases in the search area.

TABLE 4.2e

ALIAS Table	
EID	A_EID

[0330] 4.3 Object Table Decomposition

[0331] The Object table contains the following columns:

TABLE 4.3a

Object Table					
EID	AID	VID	Disting	ValueNorm	ValueRaw

[0332] The Object Table essentially contains information for finding a particular value (e.g. AID=surname, Value-Norm=“HARVEY”) and for retrieving values (e.g. AID=surname, ValueRaw=“Harvey”). This table can therefore be split into two tables: SEARCH and ENTRY.

[0333] The Search Table is used to resolve filters in the Search service. It is also used to find values during Compare, Modify and ModifyRDN. The Search table contains one row for each attribute value of each entry. Only the normalised values are stored in this table.

TABLE 4.3b

SEARCH Table				
EID	AID	VID	DISTING	NORM

[0334] The Entry table is used to return values in Reads and Searches. The Entry table contains one row for each attribute value for each entry. The RAW value is the value

exactly as initially supplied when the entry was added or modified.

TABLE 4.3c

ENTRY Table			
EID	AID	VID	RAW

[0335] 4.4 Attribute Table

[0336] The Attribute table is essentially the same as the Conceptual Design. In practice the “type” field is only descriptive, since any incoming/outgoing X.500 Object Identifier gets converted to/from the internal attribute identifier, AID. Thus this column has been renamed DESC to signify that it is a description field.

TABLE 4.4

ATTR Table			
AID	SYX	DESC	ObjectId

[0337] 4.5 Index Selection

[0338] Performance when using SQL is achieved because the RDBMS is able to satisfy the query using a relevant index. This means that every query that has a condition (the “where” clause in SQL) is preferred to have an associated index (otherwise the RDBMS has to resort to a table level scan). However in practical RDMS’s:

[0339] The number of indexes is restricted;

[0340] There may be a high overhead to maintain secondary indexes;

[0341] Composite indexes may be required to satisfy any one query. Thus, if performing a query across columns (e.g., type=surname and value “SMITH”) then separate indexes on type and value may not result in a fully indexed access. A composite index on both type and value may be required.

[0342] One innovation of the table decomposition in the previous sections is to maximise the use of primary indexes across tables. This reduces the number of secondary indexes (i.e. they become primary indexes on their own table). Following is a list of the indexes for each of the six tables used in the logical design.

TABLE 4.5

Table indexes for the Logical Design		
Table	Primary Key	Secondary Index
DIT	PARENT, RDN	EID
NAME	EID	
TREE	PATH	EID
SEARCH	AID, NORM	EID, AID, VID
ENTRY	HID, AID, VID	
ATTR	(cached)	

[0343] The table design means that many queries can be handled without joins, giving substantial performance improvement.

[0344] The joins that are considered necessary are listed below:

- [0345] List—for returning the RAW-RDNs under a given object (DIT joined with NAME).
- [0346] Search/Subtree—for finding EIDs that match a filter over a whole subtree (where the base object is not the root) (TREE joined with SEARCH).
- [0347] Search/OneLevel—for finding EIDs that match a filter one-level under the base object (DIT joined with SEARCH).
- [0348] Search/Aliases/Subtree—for finding all the aliases in a subtree (TREE joined with ALIAS).
- [0349] Search/Aliases/OneLevel—for finding all the aliases under a given object (DIT joined with ALIAS).

[0350] Note that the above joins are first level joins (i.e. between only two tables). It is preferable not to use higher order joins.

[0351] 4.6 Input/Output Performance

[0352] An innovation of decomposing tables around services, which increases the number of tables, is that the new tables are much smaller than the unfragmented tables. This can significantly reduce the amount of I/O for the following reasons:

[0353] Row Size

[0354] By reducing the number of columns in any row, the row width will be shortened. This means that more rows will fit onto a page (where it is assumed that one disk I/O returns one “page” of information). In combination with clustering below, whenever a set of rows need to be retrieved, only one (or a few) page(s) may actually have to be read off the disk (e.g. when reading the attributes of an object, if the ENTRY table is keyed on EID, AID, VID then all the rows relating to that object will be together and will probably be on the same page).

[0355] Clustering

[0356] Each of the fragmented tables is preferred to have their own (independent) primary key which enables them to cluster data according to how it is used. The primary key may dictate the “storage structure”. Thus in the SEARCH table, if the primary key is on AID, NORM (i.e. type, value) then all the data of the same type (e.g. surname) and similar values (e.g. Harvey, Harrison) will be clustered in the same area of the disk. This means that during a Search (e.g. surnames beginning with “HAR”) similar data will collected together on the one (or just a few) disk page(s). If the rows are small then the number of disk pages that have to be accessed is significantly reduced.

[0357] Caching

[0358] Most commercial RDBMS’s have the ability to cache pages frequently accessed. Since tables are effectively input (e.g. Navigating using the DIT table), or output (e.g. retrieving information from the ENTRY table) then similar requests (e.g. Searches over the same portion of the Tree) will tend to result in frequently used pages being cached, meaning frequently invoked queries will gain significant

benefits. Also the caching is more efficient since pages are “information intensive” as a result of small row size and clustering.

[0359] Management

[0360] Smaller tables are generally easier to manage: e.g. viewing, creating indexes, collecting statistics, auditing, backups, etc.

[0361] 5. Logical Methods

[0362] This section describes methods of interrogating the Logical Design tables, with reference to FIG. 2B.

[0363] Throughout this section, each X.500 method is defined and illustrated with an example. Referring again to FIG. 4, which will be referred to in the following discussion as Table 5a, it can be seen that Table 5a displays a small hierarchy tree which includes an alias reference. The corresponding Table contents are shown in Table 5b.

DIT			
EID	PARENT	ALIAS	RDN
1	0	0	DATA CRAFT
10	1	1	NETWORKS
11	1	0	SALES
12	1	0	MARKETING
20	11	0	NETWORK PRODUCTS
30	20	0	CHRIS MASTERS
31	20	0	ALANA MORGAN
32	20	0	PETER EVANS

[0364]

NAME	
EID	RAW
1	[Datacraft]
10	[Networks]
11	[Sales]
12	[Marketing]
20	[Network Products]
30	[Chris Masters]
31	[Alana Morgan]
32	[Peter Evans]

NOTE: [...] indicates a binary encoding of the exact data entry value.

[0365]

TREE	
BID	PATH
1	1.
10	1.10.
11	1.11.
12	1.12.
20	1.11.20.
30	1.11.20.30.
31	1.11.20.31.
32	1.11.20.32.

[0366]

ALIAS	
ED	A-EID
10	20

[0367]

ATTRIBUTE			
AID	SYX	DESC	OBJECTID
0	objectIdentifierSyntax	objectClass	2.5.4.0
1	distinguishedNameSyntax	aliasedObjectName	2.5.4.1
3	caseIgnoreStringSyntax	commonName	2.5.4.3
4	caseIgnoreStringSyntax	surname	2.5.4.4
7	caseIgnoreStringSyntax	localityName	2.5.4.7
8	caseIgnoreStringSyntax	stateOrProvinceName	2.5.4.8
9	caseIgnoreStringSyntax	streetAddress	2.5.4.9
10	caseIgnoreStringSyntax	organizationName	2.5.4.10
11	caseIgnoreStringSyntax	organizational-UnitName	2.5.4.11
12	caseIgnoreStringSyntax	title	2.5.4.12
13	caseIgnoreStringSyntax	description	2.5.4.13
16	PostalAddress	postalAddress	2.5.4.16
17	caseIgnoreStringSyntax	postalCode	2.5.4.17
18	caseIgnoreStringSyntax	postOfficeBox	2.5.4.18
20	telephoneNumberSyntax	telephoneNumber	2.5.4.20

[0368]

SEARCH				
EID	AID	VID	DISTING	NORM
1	0	0	0	2.5.6.4
1	10	0	1	DATAACRAFT
1	16	0	0	266-268 MAROONDAH HIGHWAY
1	17	0	0	3138
10	0	0	0	2.5.6.1
10	1	0	1	DATAACRAFT/SALES/NETWORK PRODUCTS
11	0	0	0	2.5.6.5
11	11	0	1	SALES
11	13	0	0	SALES DEPARTMENT
12	0	0	0	2.5.6.5
12	11	0	1	MARKETING
12	13	0	0	MARKETING DEPARTMENT
20	0	0	0	2.5.6.5
20	11	0	1	NETWORK PRODUCTS
20	13	0	0	NETWORK PRODUCTS SECTION
30	0	0	0	2.5.6.7
30	3	0	1	CHRIS
30	4	0	1	MASTERS
30	12	0	0	SALES MANAGER
30	20	0	0	03 727 9456
30	20	1	0	018 042 671
31	0	0	0	2.5.6.7
31	3	0	1	ALANA
31	4	0	1	MORGAN
31	12	0	0	SALES SUPPORT
31	20	0	0	03 727 9455
32	0	0	0	2.5.6.7
32	3	0	1	PETER
32	4	0	1	EVANS

-continued

SEARCH				
EID	AID	VID	DISTING	NORM
32	12	0	0	SALESPERSON
32	20	0	0	03 727 9454

[0369]

TABLE 5b

Example Tables ENTRY			
EID	AID	VID	RAW
1	0	0	[2.5.6.4]
1	10	0	[Datacraft]
1	16	0	[266-268 Maroondah Highway]
1	17	0	[3138]
10	0	0	[2.5.6.1]
10	1	0	[Datacraft/Sales/Network Products]
11	0	0	[2.5.6.5]
11	11	0	[Sales]
11	13	0	[SalesDepartment]
12	0	0	[2.5.6.5]
12	11	0	[Marketing]
12	13	0	[Marketing Department]
20	0	0	[2.5.6.5]
20	11	0	[Network Products]
20	13	0	[Network Products Section]
30	0	0	[2.5.6.7]
30	3	0	[Chris]
30	4	0	[Masters]
30	12	0	[Sales Manager]
30	20	0	[(03) 727-9456]
30	20	1	[(018)-042 671]
31	0	0	[2.5.6.7]
31	3	0	[Alana]
31	4	0	[Morgan]
31	12	0	[Sales Support]
31	20	0	[(03) 727-9455]
32	0	0	[2.5.6.7]
32	3	0	[Peter]
32	4	0	[Evans]
32	12	0	[Salesperson]
32	20	0	[(03) 727-9454]

NOTE: [...] indicates a binary encoding of the exact data entry value.

[0370] 5.1 Common Services

[0371] Tree Navigation

[0372] All X.500 services rely on navigating the directory tree, illustrated in FIG. 3. The purpose of tree navigation is to retrieve the EID of the entry corresponding to the supplied Distinguished Name. Navigation begins from the root of the tree and continues down the tree until all the RDN's in a DN have been resolved (verified). This process is known as a "Tree Walk".

[0373] The DIT Table is the primary table used for tree navigation. Referring to the example hierarchy tree, illustrated as table 5a in FIG. 3, resolution of the DN "Datacraft/Sales/Network Products/Peter Evans" involves the following processes:

[0374] Scan the DIT table for a row containing PARENT=0 and RDN="DATAACRAFT". The EID for this row is 1.

[0375] Scan the DIT table for a row containing PAR-ENT=1 and RDN="SALES". The EID for this row is 11.

[0376] Scan the DIT table for a row containing PAR-ENT=11 and RDN="NETWORK PRODUCTS". The EID for this row is 20.

[0377] Scan the DIT table for a row containing PAR-ENT=20 and RDN="PETER EVANS". The EID for this row is 32.

[0378] The DN has now been resolved and any values relating to the object can be obtained from the Entry Table using the key EID=32.

[0379] Aliases

[0380] Sometimes a DN can contain an alias, which is effectively another DN. Aliases complicate the tree walk process because the tree walk cannot continue until the alias is resolved. This requires a separate tree walk for the alias.

[0381] As an example, consider the DN "Datacraft/Networks/Peter Evans". The first two steps in resolving this DN would be:

[0382] Scan the DIT table for a row containing PAR-ENT=0 and RDN="DATACRAFT". The EID for this row is 1.

[0383] Scan the DIT table for a row containing PAR-ENT=1 and RDN="Networks"

[0384] The EID for this row is 10.

[0385] At this stage we discover that this entry is an alias. The Alias Table is checked to see if the EID of the alias has been cached. If this is the first time an attempt has been made to resolve this alias then the A_EID column in the Alias Table will be zero. For the purpose of discussion it will be assumed that this is the first time.

[0386] To resolve the alias, the DN of the aliased object must be determined. This is stored in the "aliasedObjectName" attribute of the alias entry. The aliasedObjectName has an AID=1 (from the ATTR table) and so the DN is obtained from the Entry Table (RAW value) where EID=10 and AID=1.

[0387] In this example, the DN of the alias is "Datacraft/Sales/Network Products". This DN is resolved completely using the normal tree walking technique. The value of EID is 20.

[0388] At this stage, navigation continues for the unresolved RDN's in the original DN, namely "PETER EVANS". The last step required is then:

[0389] Scan the DIT table for a row containing PAR-ENT=20 and RDN="PETER EVANS".

[0390] Once an alias has been resolved it can be added (cached) in the Alias Table. This table contains a reference, A_EID, to the aliased object. In the above example, an entry in the Alias Table with an EID of 10 would have an A_EID of Once an alias has been cached a tree walk is no longer necessary to resolve the alias.

[0391] Directory Paths

[0392] When objects are added to the DIT table, a corresponding row is added to another table called the Tree Table. This table stores the list of the EID's which identify a "Path" to the object.

[0393] Distinguished Names

[0394] Most services require the distinguished name to be returned in the Service Result. Using the directory path from the Tree Table, a DN can be constructed from the RAW RDN values stored in the Name Table.

[0395] Entry Information Selection

[0396] Many of the X.500 Services are requested with an argument called "EntryInformationSelection" or EIS. The EIS argument is used to indicate what information in the Entry should be returned. Basically, EIS can be optionally;

[0397] no information

[0398] attributes and values for selected or all attributes

[0399] values only for selected or all attributes

[0400] Entry Information

[0401] Entry Information is a return parameter for Read and Search. It always contains the Distinguished Names of selected entries and, optionally, attributes and/or values as specified in the EIS argument of the request.

[0402] Common Arguments

[0403] All of the X.500 Services pass a set of common arguments in the Service Request. Common Arguments contain information such as service controls (time limit and size limit), the DN of the requestor of the service and security information.

[0404] Common Results

[0405] Some X.500 Services pass a set of common results in the Service Response. Common Results contain information such as security parameters, the DN of the performer of the service and an alias dereferenced flag.

[0406] 5.2 Read Service

[0407] A Read operation is used to extract information from an explicitly identified entry.

X.500 definition	
Argument Name	Description
EntryInformationSelection	A Distinguished Name The attributes and values to be returned (ie EIS)
Common Arguments Result	Description
Entry Information	The DN plus any attributes and values returned
Common Results	

[0408] Method

[0409] Perform a tree walk using the DIT table, resolving aliases if necessary. Obtain the base EID.

[0410] Using PATH from the Tree Table and the RAW RDN's from the Name Table, build a DN.

- [0411] If EIS specifies no attributes or values, just return the DN.
- [0412] If EIS specifies ALL types and values, return the RAW values from the Entry Table for the matching EID.
- [0413] If EIS specifies selected types and values, obtain the AID's from the Attribute Table and then return selected types and/or values for the matching EID.

EXAMPLE

- [0414] Read the entry "Datacraft/Sales/Network Products/Peter Evans".
- [0415] EIS is set to: attribute Types=allAttributes, InfoTypes=attributeTypesAndValues.
- [0416] Using the DIT table perform a Tree Walk traversing EID's 1, 11, 20 and 32 for the normalised RDN's DATA-CRAFT, SALES, NETWORK PRODUCTS, PETER EVANS. The EID of the selected object is 32.
- [0417] Extract the PATH from the Tree Table for EID=32. The PATH is 1.11.20.32.
- [0418] Build aDN from the RAW values in the Name Table for EID's 1, 11, 20, 32.
- [0419] Using the Entry Table and the Attribute Table, for each matching EID;
 - [0420] return the OBJECTID's from the Attribute Table and the ASN.1 encoded RAW values from the Entry Table

2.5.4.0	[2.5.6.7]
2.5.4.3	[PETER]
2.5.4.4	[EVANS]
2.5.4.9	[SALESPERSON]
2.5.4.20	[(03) 727-9454]

- [0421] return the DN 5.3 Compare Service

[0422] A Compare operation is used to compare a value (which is supplied as an argument of the request) with the value(s) of/particular attribute type in a particular object entry.

X.500 DEFINITION	
Argument	Description
Name	A Distinguished Name
AttributeValueAssertion	The attribute type and value to be compared
Common Arguments	
Result	Description
DistinguishedName	The DN of the selected object (returned if an alias is dereferenced)
matched	TRUE / FALSE result of compare
fromEntry	N/A
Common Results	

- [0423] Method

[0424] Perform a tree walk using the DIT table, resolving aliases if necessary. Obtain the EID of the base object.

- [0425] From the Attribute Table, obtain the AID of the attribute to be compared.
- [0426] From the Entry Table, select the row(s) matching the EID and AID.
- [0427] Compare the value.
- [0428] Return TRUE or FALSE as the Compare result.
- [0429] If an alias is dereferenced, return the DN of the selected object, using the path from the Tree Table and the RAW RDN's from the Name Table.

EXAMPLE

- [0430] Compare the DN "Datacraft/Sales/Network Products/Peter Evans" with a purported AttributeValueAssertion of "title=[Salesperson]".
- [0431] Obtain the EID for the given DN using a TreeWalk. The EID of the selected object is 32.
- [0432] Using the Attribute table, obtain the AID for "title", ie AID=12.
- [0433] Using the Search Table locate rows with EID=32 and AID=12 and test for "NORM=SALESPERSON".
- [0434] Return TRUE or FALSE depending on the outcome of this test. In this instance the result would be TRUE.
- [0435] Since no aliases were dereferenced, the DN of the entry is not returned.
- [0436] 5.4 List Service
- [0437] A list operation is used to obtain a list of immediate subordinates of an explicitly identified entry.

X.500 DEFINITIONz	
Argument	Description
Name	A Distinguished Name
Common Arguments	
Result	Description
DistinguishedName	The DN of the selected object (returned if an alias is dereferenced)
subordinates	A list of RDN's for the subordinate entries (aliases, indicated by an alias flag, are not dereferenced)
partialOutcomeQualifier	An indication that an incomplete result was returned, eg, a time limit or size limit restriction.
Common Results	

- [0438] Method

- [0439] Perform a tree walk using the DIT table, resolving aliases if necessary. Obtain the EID of the base object.
- [0440] Using the DIT and Name Tables return the ALIAS flag and the RAW RDN PARENT is equal to the EID of the base object.

EXAMPLE

- [0441] Perform a list for the DN "Datacraft".
- [0442] Obtain the EID for the DN using a TreeWalk. The EID of the selected object is "1".

[0443] For each EID with a PARENT=1

[0444] return the RAW RDN from the Name Table, ie, [Networks], [Sales], [Marketing]

[0445] return the alias flags, ie, TRUE, FALSE, FALSE.

[0446] As no alias was dereferenced in the tree walk, the DN of the selected object is not returned. Note also that the alias entry [Networks] is not dereferenced.

[0447] 5.5 Search Service

[0448] The Search Service is the most complex of all X.500 services. Search arguments indicate where to start the search (baseObject), the scope of the search (subset), the conditions to apply (filter) and what information should be returned (selection). In addition, a flag is passed to indicate whether aliases should be dereferenced (searchAliases).

[0449] The possible values for subset are baseObject, oneLevel and wholesubtree. Base object indicates that the search filter will only be applied to attributes and values within the base object. OneLevel indicates the Search filter will be applied to the immediate subordinates of the base object. Whole subtree indicates the Search filter will be applied to the base object and all of its subordinates.

[0450] A simple example of a filter condition would be: surname="EVANS" or telephoneNumber PRESENT.

X.500 DEFINITION	
Argument	Description
baseObject	The Distinguished Name of the baseObject
subset	baseObject, oneLevel or wholeSubtree
filter	search conditions
searchAliases	a flag to indicate whether aliases among subordinates of the base object should be dereferenced during the search.
selection	EIS as for READ. The attributes and values to be returned.
Common Arguments	
Result	Description
DistinguishedName	The DN of the selected object (returned if an alias is dereferenced)
entries	Attributes & values (as defined in selection) for the entries which satisfy the filter.
partialOutcomeQualifier	An indication that an incomplete result was returned, eg, a time limit or size limit restriction.
Common Results	

[0451] The search procedures for each search scope are outlined as follows:

[0452] Base Object

[0453] Perform a tree walk using the DIT table, resolving aliases if necessary. Obtain the EID of the base object.

[0454] Apply the filter to attributes and values in the Search Table with the EID of the selected object.

[0455] If the filter condition is matched, return the Entry Information from the Entry Table.

[0456] If an alias is dereferenced, return the DN using the Tree Table to extract the PATH and the Name Table to build the DN.

[0457] One Level

[0458] Perform a tree walk using the DIT table, resolving aliases if necessary. Obtain the EID of the base object.

[0459] Check to see if any aliases exist with PARENT=EID and if so resolve them to obtain an aliases dereferenced list.

[0460] Using the Search and DIT Tables, apply the filter (attribute/value conditions) and the scope (PARENT=EID of selected object and any aliases dereferenced). A list of matching EID's will be returned.

[0461] If an alias is dereferenced, return the DN using the Tree Table to extract the PATH and the Name Table to build the DN.

[0462] For each matching EID:

[0463] Return the Entry Information obtained from the Search Table using the Entry Table (as per Read Service).

[0464] Whole Subtree

[0465] Perform a tree walk using the DIT table, resolving aliases if necessary. Obtain the EID of the base object.

[0466] Check to see if any aliases exist with PATH prefix matching the PATH of the selected object.

[0467] For each alias discovered, check to see if the alias points outside the current subtree and if it does repeat the previous step. Once all aliases have been resolved, a set of unique base objects will have been found (with no overlapping areas).

[0468] Using the Search and Tree Tables, apply the filter (attribute/value conditions) and the scope (PATH LIKE PATH prefix of the selected object) to each unique base object. A list of matching EID's will be returned.

[0469] If an alias is dereferenced during Navigation (not during searching), return the DN using the Tree Table to extract the PATH and the Name Table to build the DN.

[0470] For each matching EID:

[0471] Return the Entry Information obtained from the Search Table using the Entry Table (as per Read Service).

EXAMPLE

[0472] Perform a search on the baseObject "Datacraft/Sales" with:

[0473] Scope set to WholeSubtree

[0474] a Filter of "surname, substring initial=M". (Look for all surnames beginning with "M")

[0475] SearchAliases set to TRUE.

[0476] EIS set to attribute Types=allAttributes, InfoTypes=attributeTypesAndValues.

[0477] Method

[0478] Obtain the EID for the base object DN using a TreeWalk. The EID of the base object is “11”.

[0479] From the Tree Table, obtain the PATH for EID=11, ie, “1.11”.

[0480] Check for any aliases among entries that have a path beginning with “1.11.”. There are no aliases in this case.

[0481] Obtain the AID for the attribute “surname” in the Attribute Table, ie, 4.

[0482] Apply the filter and scope simultaneously. i.e. Using the Search Table, obtain a list of EID’s from the target list where AID=4 and the value begins with “M” joined with the Tree Table who’s PATH is LIKE ‘1.11.%’. The matching EID’s are 30 and 31.

[0483] Using the Entry Table and the Attribute Table, for each matching EID:

[0484] return the OBJECTID’s from the Attribute Table and the ASN.1 encoded RAW values from the Entry Table

i.e.,	2.5.4.0,	[2.5.6.7],
	2.5.4.3,	[Chris],
	2.5.4.4	[Masters]
	2.5.4.9	[Sales Manager]
	2.5.4.20	[(03) 727-9456]
	2.5.4.20	[(018) - 042 671]
	2.5.4.0	[2.5.6.7]
	2.5.4.3	[Alana]
	2.5.4.4	[Morgan]
	2.5.4.9	[Sales Support]
	2.5.4.20	[(03) 727-9454]

[0485] 5.6 Add Entry Service

[0486] An AddEntry operation is used to add a leaf entry either an object entry or an alias entry) to the Directory Information Tree.

X.500 DEFINITION	
Argument	Description
object	The Distinguished Name of the entry to be added
entry	A set of attributes to add
Common Arguments	
Result	Description
NULL	NULL

[0487] Method

[0488] Using the DIT table, tree walk to the parent of the entry to be added (Parent EID).

[0489] Using the DIT table, check if the entry exists (check for RDN=new RDN and PARENT=Parent EID).

[0490] If the entry does not exist, allocate a new EID and add the entry. Insert into the DIT Table, the Name Table, the Tree Table, the Search Table, the Entry Table and, if it is an alias entry, the Alias Table.

EXAMPLE

[0491] Under the object with a DN of “Datacraft/Marketing” add an object with the following attributes and values.

surname	[Delahunty]
commonName	[Mary]
title	[Marketing Manager]
telephoneNumber	[(03) 727-9523]

[0492] Obtain the EID for the base object DN using a TreeWalk. The EID of the base object is “12”.

[0493] Using the DIT Table, look for a duplicate entry, ie, PARENT=12 and RDN=“MARY DELAHUNTY”. No duplicates exist.

[0494] Add the following rows to the Tables shown.

DIT			
EID	PARENT	ALIAS	RDN
33	11	0	MARY DELAHUNTY

[0495]

NAME	
EID	RAW
33	[Mary Delahunty]

[0496]

TREE	
EID	PATH
33	1.12.21.

[0497]

SEARCH				
EID	AID	VID	DISTING	NORM
33	0	0	0	2.5.6.7
33	3	0	1	DELAHUNTY
33	4	0	1	MARY
33	12	0	0	MARKETING
				MANAGER
33	20	0	0	03 727 9523

[0498]

ENTRY			
EID	AID	VID	RAW
33	0	0	[2.5.6.7]
33	3	0	[Delahunty]
33	4	0	[Mary]
33	12	0	[Marketing Manager]
33	20	0	[(03) 727-9523]

[0499] 5.7 Remove Entry Service

[0500] A RemoveEntry operation is used to remove a leaf entry (either an object entry or an alias entry) from the Directory Information Tree.

X.500 DEFINITION	
Argument	Description
object	The Distinguished Name of the entry to be deleted
Common Arguments	
Result	Description
NULL	NULL

[0501] Method

[0502] Perform a tree walk using the DIT table. Obtain the EID of the base object.

[0503] If the entry exists, and it is a leaf entry, then for the condition EID=EID of the selected object, delete from the DIT Table, the Name Table, the Tree Table, the Search Table, the Entry Table and, if it is an alias entry, the Alias Table.

EXAMPLE

[0504] Delete the object with a DN of “Datacraft/Marketing/Mary Delahunty”

[0505] Method

[0506] Obtain the EID for the base object DN using a TreeWalk. The EID of the base object is “21”. Check that no entries have PARENT=21.

[0507] Delete all rows added to the DIT Table, the Name Table, the Tree Table, the Search Table and the Entry Table (refer to Add Entry example) where EID=21.

[0508] 5.8 Modify Entry Service

[0509] The ModifyEntry operation is used to perform a series of one or more of the following modifications to a single entry:

- [0510] add a new attribute
- [0511] remove an attribute
- [0512] add attribute values
- [0513] remove attribute values

[0514] replace attribute values

[0515] modify an alias

X.500 DEFINITION	
Argument	Description
object	The Distinguished Name of the entry to be modified
changes	A list of modifications
Common Arguments	
Result	Description
NULL	NULL

[0516] Method

[0517] Perform a tree walk using the DIT table. Obtain the EID of the selected object.

[0518] For the selected object, perform one or more of the following actions: Add Value, Delete Value, Add Attribute, Delete Attribute.

[0519] The operations required for each action are as follows:

[0520] Add Value

[0521] If the attribute exists, add the value to the Entry Table and the Search Table. Checks are: If the attribute is single valued test for an existing value; if the attribute is multi-valued check for a duplicate value.

[0522] Delete Value

[0523] For the Entry Table and the Search Table, if the value exists, delete it. A Distinguished Value cannot be deleted.

[0524] Add Attribute

[0525] If the attribute does not exist, add the Attribute Values to the Entry Table and the Search Table.

[0526] Delete Attribute

[0527] For the Entry Table and the Search Table, if the attribute exists, delete it. Delete all values with AID=attr and EID=base object. Naming attributes cannot be deleted.

EXAMPLE

[0528] Modify the Entry “Datacraft/Sales/Network Products/Chris Masters” with the following changes:

- [0529] Delete Attribute and Value telephoneNumber 018-042 671
- [0530] Modify Attribute and Value title Sales Assistant
- [0531] The Search and Entry Tables reflect the changes.

SEARCH				
EID	AID	VID	DISTING	NORM
30	0	0	0	2.5.6.7
30	3	0	1	CHRIS
30	4	0	1	MASTERS
30	12	0	0	SALES ASSISTANT
30	20	0	0	03 727 9456

[0532]

ENTRY			
EID	AID	VID	RAW
30	0	0	[2.5.6.7]
30	3	0	[Chris]
30	4	0	[Masters]
30	12	0	[Sales Assistant]
30	20	0	[(03) 727-9456]

[0533] 5.9 Modify RDN Service

[0534] The ModifyRDN operation is used to change the Relative Distinguished Name of a leaf entry (either an object entry or an alias entry) from the Directory Information Tree.

Arguments	Description
object	The Distinguished Name of the entry to be modified
newRDN	The new RDN of the entry
deleteOldRDN	flag - delete all values in the old RDN not in new RDN
Common Arguments	
Result	Description
NULL	NULL

[0535] Method

- [0536] Perform a tree walk using the DIT table. Obtain the EID and Parent EID of the base object.
- [0537] Using the DIT table, check for equivalent entries and return error if one is found. An equivalent entry has RDN=new RDN and PARENT Parent EID.
- [0538] Using the Name Table, replace the old RDN with the new RDN.
- [0539] Using the DIT Table, replace the old RDN with the new RDN.
- [0540] Using the Entry Table, insert the new value.
- [0541] Using the Search Table, locate value=old RDN and set DISTING to 0. Insert the new value.
- [0542] If deleteOldRDN is set to TRUE the procedures following the Tree Walk are as follows:
- [0543] Using the DIT table, check for a sibling with the same name and an EID not equal to the base EID

[0544] Using the Name Table, replace the old RDN with the new RDN.

[0545] Using the DIT Table, replace the old RDN with the new RDN.

[0546] Using the Entry Table, delete the old value(s) and insert the new value(s).

[0547] Using the Search Table, delete the old value(s) and insert the new value(s).

EXAMPLE

[0548] Modify the RDN of “Datacraft/Sales/Network Products/Chris Masters”. The new RDN is “Christine Masters”.

[0549] deleteOldRDN is set to FALSE.

[0550] The changes to the Tables will be as follows:

DIT			
EID	PARENT	ALIAS	RDN
21	11	0	CHRISTINE MASTERS

[0551]

NAME	
EID	RAW
21	[Christine Masters]

[0552]

SEARCH				
EID	AID	VID	DISTING	NORM
30	0	0	0	2.5.6.7
30	3	0	1	CHRISTINE
30	3	1	0	CHRIS
30	4	0	1	MASTERS
30	12	0	0	SALES ASSISTANT
30	20	0	0	03 727 9456

[0553]

ENTRY			
EID	AID	VID	RAW
30	0	0	[2.5.6.7]
30	3	0	[Christine]
30	3	1	[Chris]

-continued			
ENTRY			
EID	AID	VID	RAW
30	4	0	[Masters]
30	12	0	[Sales Assistant]
30	20	0	[(03) 727-9456]

[0554] 5.10 Complications

[0555] If error, limit or abandon occurs during processing of any of the services, then the processing is discontinued and an appropriate error message returned.

[0556] Errors

[0557] Each X.500 service consists of 3 parts; ARGUMENT, RESULT and ERRORS. In the above descriptions of the services, ARGUMENT and RESULT have been included in the X.500 definitions. Error conditions, however, are many and varied and no attempt is made to describe them in this document. The National Institute of Standards and Technology (NIST) document “Stable Implementation Agreements for Open Systems Interconnection Protocols: Version 3” provides a full coverage of errors for the X.500 standard.

[0558] Time Limit & Size Limit

[0559] Time Limit and Size Limit form part of Service Controls. They can be optionally set to some finite limit and included in the Common Arguments.

[0560] Time Limit indicates the maximum elapsed time, in seconds, within which the service shall be provided. Size Limit (only applicable to List and Search) indicates the maximum number of objects to be returned. If either limit is reached an error is reported. For a limit reached on a List or a Search, the result is an arbitrary selection of the accumulated results.

[0561] Abandon

[0562] Operations that interrogate the Directory, ie Read, Compare, List and Search, may be abandoned using the Abandon operation if the user is no longer interested in the results.

[0563] Aliases & Search

[0564] If an alias is encountered in a search and that alias points to a separate branch of the directory tree, then dereferencing of the alias requires:

[0565] Navigation from the root entry to the referenced entry

[0566] Searching of all items subordinate to the referenced entry

[0567] In the example shown in FIG. 5, if a WholeSubtree Search was performed on a base object of “Telco/Corporate/Data Services” the entries “Mervyn Purvis” and the alias “Strategic” would be searched. Strategic, however, points to

a different branch of the tree which requires searching of the entry “Strategic” and all of its subordinates, ie, “Alan Bond”, “Rex Hunt”, “Wayne Carey” and “John Longmire”.

[0568] 5.11 Implementation Optimisations

[0569] The Logical methods include a number of optimisations that enhance performance. These methods are outlined below.

[0570] Caching

[0571] The Attribute table can be cached. This means that (apart from initial loading of the attributes) no SQL statements need to be issued to the database when decoding or encoding the attributes. In the present X.500 system attribute conversions are performed in memory. This provides a substantial speed advantage.

[0572] Validation

[0573] Query validation is performed in memory where possible. This avoids database rollbacks which are time and system consuming. For example when adding an entry each attribute is validated before any attempt is made to add the entry. If an error is found then no SQL calls need to be issued.

[0574] Optimise Query Handling

[0575] As the format of most services is known, many instances of these services can be resolved using static SQL statements. More complex services, such as searches with complex filters, can be resolved using dynamic SQL. This enables arbitrarily complex searches to be performed.

[0576] Parallel Queries

[0577] Also when processing search results the present system utilises set orientation queries of SQL to avoid ‘row at a time’ processing. Thus search results may be assembled in parallel in memory.

[0578] Data Storage

[0579] The tables that store raw data store the data in ASN.1 format. This provides an efficient means of transferring data into or out of the database.

[0580] Database Techniques

[0581] Complex services can be further improved by using the query optimiser, which provides a mechanism for reducing the time spent in resolving the query. The use of a relational database also provides an efficient use of memory and enables large databases to be constructed without the need for large amounts of memory being available. Many other X.500 applications cache the entire database in memory to achieve performance. This method consumes large amounts of memory and is not scalable.

[0582] 6. Physical Design

[0583] The physical design results from a process called physical transformation of the logical design. The physical design represents a preferred realisation or embodiment of the logical design. FIG. 2B and the tables below show one form of the physical design. New columns and tables are highlighted by double borders.

TABLE 6

Physical Design									
<u>DIT</u>									
EID	PARENT	RDNKEY	RDN	FLAGS					
<u>NAME</u>									
EID	RAW	FLAGS							
<u>TREE</u>									
EID	LEV1	LEV2	LEV3	LEV4	PATH	FLAGS			
<u>INFO</u>									
MAXEID	FLAGS								
<u>ALIAS</u>									
EID	A_EID	FLAGS							
<u>SEARCH</u>									
EID	AID	VID	NORMKEY	NORM	FLAGS				
<u>ENTRY</u>									
EID	AID	VID	RAW	FLAGS					
<u>BLOB</u>									
EID	AID	VID	VFRAG	RAW	FLAGS				
<u>ATTR</u>									
AID	SYNTAX	DESC	OBJECTID				FLAGS		
<u>SENTRY</u>									
EID	AID	VID	VALUE	FLAGS					
<u>OCLASS</u>									
OCID	DESC	OBJECTID	MUSTLIST	MAYLIST	SUPERLIST	FLAGS			

[0584] The reasons for the above changes are described below.

[0585] 6.1 Efficiency

[0586] INFO Table

[0587] This table holds the highest EID value that has been used in the database. The inclusion of the INFO table enables the next EID to be obtained without any calculation of the maximum EID being performed by the database. This provides improved efficiency in adding entries to the database. More importantly the inclusion of the INFO table removes contention problems which may occur when multiple DSA's are adding entries at the same time.

[0588] Shadow Keys

[0589] Three tables have had shadow keys added. These are:

- [0590] a) The NORMKEY column in the SEARCH table.
- [0591] b) The RDNKEY column in the DIT table.
- [0592] c) The LEV1, LEV2, LEV3 and LEV4 columns in the TREE table.

[0593] Each of these shadow key columns is a shortened version of a larger column. They have been added to shorten the indexes on each table. This gives improved performance for any queries that use the indexes and it also improves disk space usage as small indexes take up less space than large indexes.

[0594] The shadow keys in the PATH table utilise the structured nature of the PATH. By being a composite key then exact matching can be used in the SQL instead of the "LIKE" operator.

[0595] e.g. WHERE LEV1=1 AND LEV2=10 AND . . .

[0596] instead of WHERE PATH LIKE '1.10.%'.

[0597] If each of the LEV columns has their own index, then a sub-tree search needs to only use the base object. e.g. LEV2=10, since all objects under entry 10 will have LEV2=10.

[0598] SENTRY Table

[0599] Some types of attribute values do not need to be normalised e.g. integer, boolean, date. Instead of storing them twice (SEARCH.NORM and ENTRY.RAW) they can

be stored just once in a hybrid table called the SENTRY table. This reduces table sizes and increases storage efficiency at the cost of having to search two tables and retrieve from two tables.

[0600] OCLASS Table

[0601] Most attributes have a wide variation in their values e.g. surnames could range from AALDERS to ZYLA with a great many different values in between. However, Object Classes (whose values are ObjectIdentifiers or OIDs) have very few values e.g. in an organisation of 10,000 people, the only object classes in the directory may be for organisation, organisationalUnit and organisationalPerson (of which many may be the latter). The OCLASS table gives a numeric descriptor to an object class called an OCID. The OCID can then be stored in the SENTRY table and a mapping done whenever an Object Class is searched or retrieved. The other LIST columns store standard object class configuration information—namely the must and may contain attributes and the inherited superclasses.

[0602] 6.2 Portability

[0603] BLOB Table

[0604] This table has been included to hold “Binary Large Objects”. The maximum size of a one row entry in the ENTRY table is limited by the length of the RAW field. This means that entries must be fragmented. Fragmented entries will occupy more than one row and so a VFRAG field must be used to denote the fragment of the entry that is being stored in a particular row.

[0605] There are two options for storing very large values:

[0606] a) Add a “fragment flag” to the ENTRY table and store the entry in fragments over a number of lines; or

[0607] b) Add a BLOB table to store the entry and add a “BLOB flag” to the ENTRY table to indicate that this value is stored in the BLOB.

[0608] The second option has a number of advantages. Firstly, the inclusion of a BLOB table prevents the ENTRY table from becoming excessively large. Generally most entries will be less than a few hundred characters in length, so the length of the RAW field in the ENTRY table can accordingly be reduced to cater for those entries and the RAW field in the BLOB table can be increased to a value approaching the maximum record size. This will make storage more efficient, i.e. reduce the amount of unused bytes in each column of each table and reduce the number of fragments needed for each entry in the BLOB table. It also means that each value will have only one entry in the ENTRY table and that the ENTRY and SEARCH tables maintain their one-to-one correlation. Secondly the use of a BLOB table enables the application to make use of any database support for Binary Large Objects. (e.g. 64 K Binary Columns).

[0609] 6.3 Functional Extensibility

[0610] FLAGS Columns

[0611] FLAGS column(s) are preferred to be added. These column(s) have been added to provide extensibility to the design. Specific values can be added to the flags as new functionality is required, without changing the table structure.

[0612] Note:

[0613] a) In the SEARCH table, the DISTING field may be absorbed into the FLAGS field.

[0614] b) In the DIT table, the ALIAS field may be absorbed into the FLAGS field.

[0615] The FLAGS column(s) may also provide a “summary” function for each of the tables. This means that the nature of an entry can be determined to some extent by checking the value of the FLAGS field. For example, a flag can be set, in the DIT table, when an entry is a leaf. Checking this flag is much simpler than checking for children of the entry.

[0616] The FLAGS column can also be used to store security information, whether an alias points inside its parents sub-tree, whether a value is a BLOB, etc.

[0617] 7. Example Implementation

[0618] The following provides an example of system performance and capabilities. It is to be understood that the present inventions should not be limited to the following disclosure.

[0619] 7.1 Overall System Benefits

[0620] The present invention is considered to provide enhanced performance over prior art implementations. Performance can be appraised in many ways, including:

[0621] aliases;

[0622] size (use of relational theory);

[0623] complexity (use of query optimiser and search method(s));

[0624] extensibility (use of meta-data); and

[0625] substantially without degrading efficiency (use of disk based model) and

[0626] reliability (use of RDBMS).

[0627] The present invention is considered unique in its ability to claim performance improvement in all areas noted above.

[0628] 7.2 Test Results

[0629] Performance testing of the present invention has been carried out, with the objectives of:

[0630] Proving that an SQL based X.500 application can perform at subsecond speeds, dispelling a widely held myth in the marketplace that it is impossible to implement an X.500 DSA application as an integrated RDBMS application and achieve efficiency and performance.

[0631] Proving that the design of an SQL based X.500 application can outperform existing memory resident style X.500 designs, especially for databases in excess of 100K entries, a typical limit of current designs.

[0632] Providing a structured suite of tests that can demonstrate the above performance on demand for a wide variety of services and database sizes.

[0633] Test results reveal the following Table 7A

TABLE 7A

Service			Database Size (number of entries)					
Operation	Qualifier	Detail	1K	10K	20K	50K	100K	200K
BIND	anonymous		0.00	0.00	0.00	0.00	0.00	0.00
LIST	level 1 4 items		0.05	0.05	0.05	0.05	0.05	0.05
	level 3 4 items		0.06	0.06	0.06	0.06	0.06	0.06
	level 4 100 items		0.22	0.23	0.23	0.24	0.23	0.24
READ	level 4 1 item, all info		0.07	0.07	0.07	0.07	0.07	0.08
	level 4 (via alias) 1 item, all info		0.07	0.07	0.07	0.07	0.07	0.07
SEARCH	1 level, equality	100 entries, 1 item	0.12	0.12	0.12	0.12	0.13	0.13
	1 level, initial	100 entries, 1 item	0.13	0.14	0.15	0.15	0.15	0.14
	1 level, any	100 entries, 1 item	0.30	0.35	0.33	0.32	0.36	0.29
	1 level, final	100 entries, 1 item	0.24	0.35	0.31	0.30	0.35	0.28
	subtree, equality	1K, 1 item, level 1	0.11	0.11	0.11	0.11	0.11	0.11
		10K, 1 item, level 1	xxx	xxx	0.12	0.12	0.12	0.12
		20K, 1 item, level 1	xxx	xxx	xxx	0.12	0.13	0.12
		50K, 1 item, level 1	xxx	xxx	xxx	xxx	0.13	0.13
		100K, 1 item, level 1	xxx	xxx	xxx	xxx	xxx	0.12
	subtree, initial	1K, 1 item, level 1	0.13	0.12	0.12	0.12	0.12	0.11
		10K, 1 item, level 1	xxx	xxx	0.11	0.12	0.12	0.12
		20K, 1 item, level 1	xxx	xxx	xxx	0.13	0.12	0.12
		50K, 1 item, level 1	xxx	xxx	xxx	xxx	0.13	0.12
		100K,1 item, level 1	xxx	xxx	xxx	xxx	xxx	0.11
	full, complex OR	all entries, 1 item	0.09	0.09	0.09	0.09	0.09	0.09
	full, complex AND	all entries, 1 item	0.11	0.11	0.11	0.11	0.11	0.11
	full, complex OR/AND	all entries, 1 item	0.26	0.28	0.29	0.28	0.29	0.26
	full, complex AND/OR	all entries, 1 item	0.12	0.12	0.13	0.14	0.13	0.12
	full, complex AND/AND	all entries, 1 item	0.16	0.15	0.16	0.17	0.18	0.18
	full, complex all entries,	1 item AND/AND/AND	0.18	0.18	0.18	0.19	0.20	0.26
	full, equality all entries,	1 item	0.08	0.08	0.08	0.08	0.08	0.08
	full, no filter, all-info	all entries, 10 items	0.30	0.74	0.43	0.59	0.49	0.67
	full, no filter, all-info	all entries, 100 items	1.36	1.84	1.50	1.79	1.82	1.86
	full, initial all entries,	1 item	0.08	0.08	0.08	0.08	0.08	0.08
ADD	level 5 100 sisters		0.22	0.19	0.22	0.20	0.19	0.19
MODIFY	level 5 100 sisters		0.09	0.11	0.11	0.11	0.11	0.11
RENAME	level 5 100 sisters		0.15	0.16	0.15	0.16	0.16	0.15
DELETE	level 5 100 sisters		0.17	0.16	0.17	0.17	0.17	0.19
UNBIND			0.00	0.00	0.00	0.00	0.00	0.00

Notes:
1. All searches and reads return all info
2. All tests were performed under the following environment; Sun SparcStation 5 with 32Mb of memory (entry level UNIX machine)
Ingres 6.4/04 configured for 32 users (standard Ingres installation)
DSA prototype V2.1.2
Timings measured at DSA console (ie does not include network overheads)
All numbers are in units of seconds and "K" means 1,000's.

[0634] 7.3 Test Conclusions

[0635] A set of directories was constructed ranging from 1K to 200K entries with varying depth and width of the hierarchy, and a corresponding test plan was produced. The tests were performed a number of times to ensure consistency.

[0636] The following conclusions can be drawn from these results;

- [0637] 1. The effects of navigation, in test, were negligible.
- [0638] 2. Reading an object via an alias, in test, showed no appreciable decrease in performance and in some cases reading an object via an alias was in

fact faster than reading the object directly. This is due to the reduced navigation required when an alias points "down" to an object that is deeper in the tree structure than the alias entry.

[0639] 3. Search results were "flat" over different sized subtrees in different sized directories for both exact and initial string searches.

[0640] 4. Initial and exact full tree searches, in test, were slightly quicker than their respective subtree searches, even though the number of entries searched was greater. This is due to the fact that the full tree searches are able to use more efficient SQL (no table joins are required).

[0641] 5. All services were, in test, performed in under one second, except for searches returning large amounts of data. However the average time of retrieval per entry drops as the number of entries retrieved increases (e.g for 10 entries retrieval time is approximately 50 milliseconds per entry, for 100 entries this drops to approximately 20 milliseconds per entry).

[0642] 6. All complex searches, in test, were performed in under one second. However, there may be some obscure searches (e.g containing combinations of NOT) which may not perform as well.

[0643] Because this is a disk based system (rather than a memory based system) performance is essentially only dependent on the number of entries actually returned. It is relatively independent of the search complexity, the depth of the hierarchy, the number of attributes per entry or the types of attributes used in the query. In a "live" application of the system it may be possible to improve on the achieved test results by tuning the caching parameters, and by having a greater diversity of attributes.

We claim:

1. A directory service arrangement comprising a data base using a plurality of tables, each table having a plurality of rows and columns, said tables being organised corresponding to their function.

2. A directory service arrangement as claimed in claim 1 wherein said plurality of tables comprises at least a first table allowing a new attribute type to be defined by adding a row to the table, a second table defining the attributes within each object, and a third table defining the relationship between objects.

3. A directory service arrangement comprising:

a database using a plurality of tables, each table having a plurality of rows and columns, and storing arbitrary data, and

means for processing said arbitrary data using a fixed set of queries/services.

4. A directory service arrangement as claimed in claim 3, wherein said database is a relational database.

5. A directory service arrangement as claimed in claim 4, wherein said database supports the structured query language (SQL).

6. A method of providing tables in a directory service apparatus, comprising:

establishing a plurality of tables in a directory service apparatus, each of said tables comprising a plurality of columns, and

defining the plurality of tables on the basis of corresponding functions related to said services.

7. A method as claimed in claim 6, wherein each service is modeled and the relationships among said services is predetermined.

8. A method as claimed in claim 6, in which services provided by the directory service are performed at least partially in a database that supports SQL.

9. A method as claimed in claim 7, in which services provided by the directory service are performed at least partially in a database that supports SQL.

10. A method of processing arbitrary data in a directory service apparatus, said apparatus comprising a plurality of

tables established on the basis of respective functions, the method including the step of:

processing the arbitrary data using a fixed set of queries/services.

11. A method of managing data in a directory service arrangement supporting at least one of X.500 or LDAP services, wherein data is identifiable by a predetermined criteria and is stored in a database medium, the method comprising the step of:

clustering data of the same criteria in the same area of the database medium.

12. A method as claimed in claim 11, wherein the criteria comprises data type.

13. A method as claimed in claim 11, wherein the criteria comprises entry identifier (EID).

14. A method as claimed in claim 11, wherein the criteria comprises a parent identifier.

15. A method as claimed in claim 11, wherein the criteria comprises data type and value.

16. A method as claimed in claim 11, wherein the criteria comprises parent identifier and entry name.

17. A method as claimed in claim 11, wherein the criteria comprises attribute type for the purpose of a search.

18. A method as claimed in claim 11, wherein the criteria comprises attribute type and similar value(s) for the purpose of a search.

19. A method as claimed in claim 11, further comprising: searching clustered data on the basis of the criteria.

20. A method as claimed in claim 12, further comprising: searching clustered data on the basis of the criteria.

21. A method as claimed in claim 13, further comprising: searching clustered data on the basis of the criteria.

22. A method as claimed in claim 14, further comprising: searching clustered data on the basis of the criteria.

23. A method as claimed in claim 15, further comprising: searching clustered data on the basis of the criteria.

24. A method as claimed in claim 16, further comprising: searching clustered data on the basis of the criteria.

25. A method as claimed in claim 17, further comprising: searching clustered data on the basis of the criteria.

26. A method as claimed in claim 18, further comprising: searching clustered data on the basis of the criteria.

27. The invention as herein disclosed.

28. A method of managing data as claimed in claim 11, wherein the step of clustering data comprises:

clustering data of similar criteria for the purpose of retrieving data.

29. A method of managing data as claimed in claim 12, wherein the step of clustering data comprises:

clustering data of similar criteria for the purpose of retrieving data.

30. A method of managing data as claimed in claim 13, wherein the step of clustering data comprises:

clustering data of similar criteria for the purpose of retrieving data.

31. A method of managing data as claimed in claim 14, wherein the step of clustering data comprises:

clustering data of similar criteria for the purpose of retrieving data.

32. A method of managing data as claimed in claim 15, wherein the step of clustering data comprises:

clustering data of similar criteria for the purpose of retrieving data.

33. A method of managing data as claimed in claim 16, wherein the step of clustering data comprises:

clustering data of similar criteria for the purpose of retrieving data.

34. A method of managing data as claimed in claim 17, wherein the step of clustering data comprises:

clustering data of similar criteria for the purpose of retrieving data.

35. A method of managing data as claimed in claim 18, wherein the step of clustering data comprises:

clustering data of similar criteria for the purpose of retrieving data.

36. A method of managing data as claimed in claim 28, further comprising:

retrieving clustered data on the basis of the criteria.

37. A method of managing data as claimed in claim 29, further comprising:

retrieving clustered data on the basis of the criteria.

38. A method of managing data as claimed in claim 30, further comprising:

retrieving clustered data on the basis of the criteria.

39. A method of managing data as claimed in claim 31, further comprising:

retrieving clustered data on the basis of the criteria.

40. A method of managing data as claimed in claim 32, further comprising:

retrieving clustered data on the basis of the criteria.

41. A method of managing data as claimed in claim 33, further comprising:

retrieving clustered data on the basis of the criteria.

42. A method of managing data as claimed in claim 34, further comprising:

retrieving clustered data on the basis of the criteria.

43. A method of managing data as claimed in claim 35, further comprising:

retrieving clustered data on the basis of the criteria.

44. A method of managing data as claimed in claim 17, wherein the step of clustering data comprises:

clustering said data around nodes for the purpose of navigating.

45. A method of managing data as claimed in claim 44, further comprising navigating in said database on the basis of said clustered data.

46. A method of managing data as claimed in claim 11, wherein the step of data clustering is effected by adding at least one index to each of a plurality of tables.

47. A method of managing data as claimed in claim 12, wherein the step of data clustering is effected by adding at least one index to each of a plurality of tables.

48. A method of managing data as claimed in claim 13, wherein the step of data clustering is effected by adding at least one index to each of a plurality of tables.

49. A method of managing data as claimed in claim 14, wherein the step of data clustering is effected by adding at least one index to each of a plurality of tables.

50. A method of managing data as claimed in claim 15, wherein the step of data clustering is effected by adding at least one index to each of a plurality of tables.

51. A method of managing data as claimed in claim 16, wherein the step of data clustering is effected by adding at least one index to each of a plurality of tables.

52. A method of managing data as claimed in claim 17, wherein the step of data clustering is effected by adding at least one index to each of a plurality of tables.

53. A method of managing data as claimed in claim 18, wherein the step of data clustering is effected by adding at least one index to each of a plurality of tables.

54. A method of managing data as claimed in claim 28, wherein the step of data clustering is effected by adding at least one index to each of a plurality of tables.

55. A method of managing data as claimed in claim 29, wherein the step of data clustering is effected by adding at least one index to each of a plurality of tables.

56. A method of managing data as claimed in claim 30, wherein the step of data clustering is effected by adding at least one index to each of a plurality of tables.

57. A method of managing data as claimed in claim 31, wherein the step of data clustering is effected by adding at least one index to each of a plurality of tables.

58. A method of managing data as claimed in claim 32, wherein the step of data clustering is effected by adding at least one index to each of a plurality of tables.

59. A method of managing data as claimed in claim 33, wherein the step of data clustering is effected by adding at least one index to each of a plurality of tables.

60. A method of managing data as claimed in claim 34, wherein the step of data clustering is effected by adding at least one index to each of a plurality of tables.

61. A method of managing data as claimed in claim 35, wherein the step of data clustering is effected by adding at least one index to each of a plurality of tables.

62. A method of arranging tables of a directory service arrangement, for providing directory services, comprising primary services and secondary services, the method comprising the steps of:

a. decomposing tables around primary service relationships, and

b. deriving secondary services via joins.

63. A method as claimed in claim 62, in which in said decomposing step, at least one of the following consideration are made:

(i) columns that have strong relationships are to be kept together, thereby avoiding unnecessary joins;

(ii) if the number of significant rows in a given column is independent of the other related columns, then that given column is considered a candidate for a separate table;

(iii) if a column is only used for locating information as an input, or only used for returning results as an output, then it is considered a candidate for its own table;

(iv) if a column is used as a key for more than one service, then it is considered a primary key and therefore a candidate for its own table; or

(v) keys are to be unique or at least strong and not repetitious.

64. A method of decomposing a principal database design, having a PROPERTY table comprising at least a column for object name and parent name, in a directory service system, the method including the step of:

applying functional decomposition to the PROPERTY table, and

deriving a conceptual design having a HIERARCHY table, an OBJECT table, and an ATTRIBUTE table.

65. A method of decomposing a conceptual design in a directory service system offering a plurality of services and being operative with a database having a plurality of tables, including a HIERARCHY table and an OBJECT table, each comprising a plurality of columns, the method including the steps of:

a. applying service decomposition to said HIERARCHY table, thereby deriving at least two table(s) comprising a DIT (Directory Information Tree) table having information required for navigation, a NAME table having information required for returning names, including raw relative distinguished name (RDN), a TREE table having information about an entry's path, and an ALIAS table having entries that are aliases: and

b. applying service decomposition to said OBJECT table, thereby deriving at least two tables, a SEARCH table for resolving filters in a Search service, said SEARCH table including one row for each attribute value of each entry, and an ENTRY table for returning values in Read and Search services, said ENTRY table having one row for each attribute value of each entry, and

c. deriving a logical design having tables defined in steps a and b above and an ATTR (ATTRIBUTE) table.

66. A method of decomposing a logical design of a directory service system into a physical design, comprising the steps of:

applying physical decomposition to the logical design, thereby deriving the physical design.

67. A directory service system for one or more objects, wherein said system comprises:

a plurality of tables according to a principal design, a conceptual design and a logical design,

wherein said principal design comprises at least one table which is a PROPERTY table.

68. A directory service system as claimed in claim 67, further comprising a plurality of tables according to a conceptual design.

69. A directory service system as claimed in claim 67, further comprising a plurality of tables according to a logical design.

70. A directory service system as claimed in claim 68, further comprising a plurality of tables according to a logical design.

71. A directory service system as claimed in claim 67, wherein the PROPERTY table has columns comprising at least one column for each of the following items:

an object name, comprising the name of said object;

a parent name, comprising the name of a parent of said object;

a type, comprising the type of a value of said object;

a syntax, comprising the form of a value of said object; and

value, comprising the contents of a value of said object.

72. A directory service system as claimed in claim 68, wherein the PROPERTY table has columns comprising at least one column for each of the following items:

an object name, comprising the name of said object;

a parent name, comprising the name of a parent of said object;

a type, comprising the type of a value of said object;

a syntax, comprising the form of a value of said object; and

value, comprising the contents of a value of said object.

73. A directory service system as claimed in claim 69, wherein the PROPERTY table has columns comprising at least one column for each of the following items:

an object name, comprising the name of said object;

a parent name, comprising the name of a parent of said object;

a type, comprising the type of a value of said object;

a syntax, comprising the form of a value of said object; and

value, comprising the contents of a value of said object.

74. A directory service system, having a plurality of tables arranged in a conceptual design, wherein each table comprises a plurality of columns, said system comprising:

a HIERARCHY table, adapted to contain information about objects, comprising their parents, their names, their relative positions in a hierarchy, and if they are aliases, for defining the structural relationship between objects;

an OBJECT table, adapted to contain information for finding a particular value and for retrieving values, for defining the attribute values within each object; and

an ATTRIBUTE table, adapted to contain information about the attributes of object(s).

75. A system as claimed in claim 74, wherein

the HIERARCHY table comprises a separate column for at least an entry ID (EID) which correlates each object with its hierarchy information, a Parent value, and a Name value,

the OBJECT table comprises a separate column for at least an EID value, an attribute ID (AID) value which correlates each value in the OBJECT table with its attribute information, a value identifier (VID) to identify values within an attribute in the OBJECT table, a Disting value to flag attribute values for naming an entry, and at least one value for each object attribute, and

the ATTRIBUTE table comprises a separate column for at least an AID value, a Type value which identifies object type, a Syntax value which identifies object form, and

an object identifier (ObjectID) for conversion between directory service identifiers and internal attribute identifiers.

76. A directory service system, comprising a plurality of tables, each having a plurality of columns, and being arranged in a logical design, said tables comprising one or more of:

- a directory information tree (DIT) table, having information for navigation;
- a TREE table, having information about the sequence of EID's down from the root;
- an ALIAS table, having information about entries that are aliases;
- a NAME table, having information required for returning names;
- a SEARCH table, having information used to resolve filters in a search service and find values during other services;
- an ENTRY table, which is used to return values in Reads and Searches, with one row being used for each attribute value of each entry; and
- an ATTRIBUTE table, which describes information about attributes of entry(s).

77. A directory services system as claimed in claim 76, wherein

- the DIT table has EID, Parent, Alias and RDN columns,
- the TREE table has EID and Path columns,
- the ALIAS table has EID and A_EID columns,
- the NME table has EID and Raw columns,
- the SEARCH table has EID, AID, VID, Disting, and Norm columns,
- the ENTRY table has EID, AID, VID and Raw columns, and
- the ATTRIBUTE table has AID, Syx, Desc and Objectid columns.

78. A directory services system comprising a plurality of tables, each having a plurality of columns, and being arranged in a physical design, said tables comprising one or more of the following:

- a directory information tree (DIT) table, having information for navigation;
- a TREE table, having information about the sequence of EID's down from the root;
- an ALIAS table, having information about entries that are aliases;
- a NAME table, having information required for returning names;
- a SEARCH table, having information used to resolve filters in a search service and find values during other services;
- an ENTRY table, which is used to return values in Reads and Searches, with one row being used for each attribute values of each entry;

an INFO table, which holds the highest EID value that has been used in the database;

a SENTRY table, which stores attribute values in raw form alone, where they do not need to be normalized;

a BLOB table, which stores binary large objects that are greater in size than that of the ENTRY table;

an ATTRIBUTE table, which stores information about attributes of entry(s); and

an OCLASS table, which provides information about object class(es) in entry(s).

79. A system as claimed in claim 38, wherein

the DIT table has EID, Parent, Rdnkey, Rdn and Flags columns,

the TREE table has EID, Lev1, Lev2, Lev3, Lev4, Path and Flags columns,

the ALIAS table has EID, A_EID, and Flags columns,

the NAME table has Eid, Raw and Flags columns,

the INFO table has Maxeid and Flags columns,

the SEARCH table has EID, AID, VID, Normkey, Norm and Flags columns,

the ENTRY table has EID, AID, VID, Raw and Flags columns,

the SENTRY table has EID, AID, VID, Value and Flags columns,

the BLOB table has EID, AID, VID, Vfrag, Raw and Flags columns,

the ATTRIBUTE table has AID, Syx, Desc, Objectid and Flags columns, and

the OCASS table has Ocid, Desc, Objectid, Mustlist, Maylist, Superlist and Flags columns.

80. An implementation of directory services in a RDBMS which supports a relational language, using service modeling, the implementation comprising:

an ATTRIBUTE table, where extensibility is provided by allowing the definition of a new attribute type by adding a row to the table;

an OBJECT table, which defines the attributes within each object; and

a HIERARCHY table which defines the relationship between the objects.

81. The implementation of claim 80 wherein said OBJECT table comprises normalized value columns and raw value columns.

82. The implementation of claim 80 wherein said HIERARCHY table comprises a normalized name column and a raw name column.

83. The implementation of claim 80 wherein the HIERARCHY table further comprises an alias column for indicating that an entry is an alias.

84. The implementation of claim 83 wherein the alias column comprises an alias and an A-EID column providing information about the destination to which the alias points.

85. The implementation of claim 83, wherein the HIERARCHY table comprises a parent column with parent ID information defining the parent entry, and a path column, the

path column containing information enabling a determination of the absolute position in a hierarchy and a determination if an entry is in a given subtree by its prefix.

86. A method of implementing a data service in a RDBMS data base which supports a relational language, having a fixed set of queries/services defined by service modeling, the method comprising:

defining an ATTRIBUTE table, where extensibility is addressed by allowing the definition of a new attribute type by adding a row to the table;

defining an OBJECT table, which defines the attributes within each object; and/or

defining a HIERARCHY table which defines the relationship between the objects; and

providing arbitrary data entries to said ATTRIBUTE table; and

processing said arbitrary data using said fixed set of queries/services.

87. The method of claim 86 further comprising the step of using the ATTRIBUTE table for incoming data to find the attribute ID (AID) from the directory service object ID for an outgoing data read from the database.

88. The method of claim 86 further comprising the step of utilising a path field to simultaneously apply an arbitrary filter over an arbitrary subtree.

89. The method of claim 87 further comprising storing in said at least one of said tables the path of each entry as a string of entry ID values.

90. The method of claim 86 further comprising:

caching the values from the ATTRIBUTE table; and

performing conversions and validations in memory.

91. The method of claim 86 further comprising:

decomposing the tables around primary service relationships and deriving secondary services via joins.

92. The method of claim 91 further comprising:

keeping together columns that have strong relationships to avoid unnecessary joins;

making a column a separate table if:

the number of significant rows in said column is independent of the other related columns; and

said column is only used for locating information (input) or only used for returning results (output); and

said column is used as a key for more than one service.

93. The method of claim 91 further comprising splitting a HIERARCHY table into a plurality of tables comprising columns for at least the id, name, tree and alias.

94. The method of claim 91 further comprising splitting the OBJECT table into a plurality of tables comprising at least a SEARCH table used to resolve filters in search service and find values during compare, rename and modify operations.

95. The method of claim 91 further comprising resolving an alias by determining the end of the alias object and caching the alias in an ALIAS table containing a reference to the alias.

96. The method of claim 95 further comprising accessing, adding, or deleting a row to a tree table, and storing a list of IDs which identify a path to an object, when objects are added to the DIT table.

97. A data arrangement for implementing directory services in a RDBMS which supports a relational language, having a fixed set of queries/services defined by service modeling, the data arrangement comprising:

an extensible ATTRIBUTE table, comprising information about attributes;

an OBJECT table, comprising data which defines the attributes within each object; and/or

a HIERARCHY table, comprising data which defines the relationship between the objects.

98. The data arrangement as set forth in claim 97 wherein said data in said ATTRIBUTE table is arranged as at least one column for each of type, syntax and description.

99. The data arrangement as set forth in claim 97 wherein said HIERARCHY table is arranged as at least a column for parent and child.

100. The data arrangement of claim 97, wherein said HIERARCHY table is arranged with data values comprising an entry identifier (EID) in a column, said EID for correlating each object with its hierarchy information.

101. The data arrangement of claim 97, wherein the OBJECT table is arranged with data values comprising an entry identifier (EID) in a column and an attribute identifier (AID) which correlates each value in the OBJECT table with its attribute information.

102. The data arrangement of claim 97, wherein the ATTRIBUTE table further comprises an object id column so that conversions between object identifiers and the internal attribute identifiers can be performed.

103. The data arrangement of claim 102, wherein said OBJECT table further comprises a value identifier (VID) to identify values within an attribute in the OBJECT table.

104. The data arrangement of claim 97 wherein the OBJECT table further comprises a column to flag distinguished values (DISTING) for naming an entry.

105. The data arrangement of claim 104 wherein distinguished values combine to form a relative distinguish name (RDN) which names an entry.

106. The data arrangement of claim 97 where both normalized data and raw data are stored in the database and the HIERARCHY and OBJECT tables include columns for normalized and raw inputs, respectively.

107. A directory service system implementing the method of any one of claims 11-46.

108. The method as recited in any one of claims 62-66, wherein said directory service arrangement comprises one of a LDAP and X.500 directory service.

109. The directory service system as recited in any one of claims 67-79 wherein said system is one of a X.500 and a LDAP service system.

110. The implementation of a directory service system as claimed in any one of claims 80-85 wherein said system is one of X.500 and a LDAP service system.

111. The method as recited in any one of claims 86-96 wherein said directory service comprises one of a LDAP and X.500 directory service.

112. The data arrangement of any one of claims **97-106**, in which the directory services are X.500 or LDAP services.

113. A computer program product comprising a computer program storage medium containing therein a computer

program operable in accordance with the method recited in any one of claims **6-66** and **86-96**.

* * * * *