Application of Novel Technologies to the Management of a Very Large Data Base

Philippe PENNY

COPERNIQUE S.A.
10-12 rue Y Tourgueneff
78380 BOUGIVAL
FRANCE

Michel PICARD

CNET PAA/OGE/TDT
38-40 rue General Leclerc
92131 ISSY LES MOULINEAUX
FRANCE

Abstract : An electronic delivery system project managing all the technical documentation related to the French telephone switching network, is presented. It is shown that several new techniques must be used to implement and to access such a large database : image acquisition and compression, optical disks management, wide band networks and high database efficiency. This paper focuses on the database related problems, especially on the very high performance required for such an application. A commercial (COPERNIQUE) database machine and an intelligent disk help in achieving this goal, and we describe their use.

1 Introduction

1.1 General context

Some years ago, the idea of electronic document delivery arose. It involves an association between a classical bibliographical database and a system storing and transmitting document images. Some projects, as ADONIS in Europe, or the National Library of Medicine's in the USA, tried to implement these ideas. However, technical and economic limits were quickly reached because users were not yet ready to pay a high cost for such a service.

Nowadays, many companies are about to pay the price in order to be able to get the right technical documentation on a damaged site of a complex system (a telephone network, a nuclear power plant, a long-range aircraft). Thus, it is reasonable to think that the first field using electronic document delivery systems will be technical documentation.

1.2 The SARDE project

CNET (the French National Center for Research in Telecommunications) is developing a project (SARDE), which is aimed at replacing the current technical document delivery process by an entirely electronic system, using a very large data base. The technical documentation concerns telephone network switching centers, and helps the technical staff in maintaining equipments which are in constant evolution. The goal of the SARDE project is an immediate access to the documents of interest.

In this system, documents are first digitalized and stored into an image database implemented on optical disks : this is the acquisition step. The document characteristics are then introduced into a reference database implemented on a database machine : this is the indexation step. Once these two steps are executed, the document is available for query.

Each query issues a command which uses special devices for database access : the DIRAM 32 programmable, intelligent disk is used to boost database searching during the first phase of the project; the DORSAL 32 database machine will later manage most of the application. DIRAM 32 and DORSAL 32 are manufactured by the French company COPERNIQUE, which is partly in charge of the analysis and the development of the reference database. Selected documents are displayed on workstations, which are powerful microcomputers, with a very high resolution screen. A 64 Kbps link is used to send images in a few seconds.

1.3 Reader's guide

The application is presented in section 2. Section 3 outlines the results of the preliminary analysis, focusing on main choices and on database architecture issues. The general architecture of the system is detailed in section 4. Section 5 presents the programmable, intelligent disk DIRAM 32. Section 6 unveils the DORSAL 32 database machine. Data security issues are addressed in section 7, and some conclusions are drawn in section 8.

2 Characteristics of the final application

2.1 Requirements

The whole database must be on-line, providing geographically distributed users with all the information related to the equipments they have to maintain. Answers to their questions must be exhaustive, data must be consistent, and response time, short. This ideal behaviour is not easy to reach because of constraints presented below. Furthermore, the system should work on (relatively) cheap equipments, upgradable without loss of performance as the volume of data and/or queries rises.

2.2 Some quantitative considerations

The documentation related to telephone switching centers represents a volume of 4 to 5 millions of pages. Half of them are made of text, and the rest comprises diagrams and drawings (average size : 19"x23"). This documentation evolves very quickly, since one third of it is renewed every year.

Up to 2000 sites must be equipped with workstations in order to query the database. The frequency and the volume of queries is not well known, but it is reasonable to predict that if it works correctly, the system will be accessed by many users simultaneously (from 100 to 200 during rush hours).

2.3 Database related problems

Commercially available database management systems (DBMS) could not be used to implement this application, because of characteristics related to the volume of data, short response time and upgradability required. One could believe that, in order to get an electronic delivery system, it is enough to add an image management system to a preexisting reference database. Our case shows that this is not always possible. Current reference DBMSs run on large - mainframes, which are not of interest for the SARDE project. It is possible to implement the same database level of service as theirs on much cheaper, dedicated equipments, i. e. one DORSAL 32 database machine connected to a mini-computer. This architecture can be upgraded at low cost, and it provides a high throughput fitting our volume and time requirements.

In a large-scale electronic delivery system, the question of performance is essential. The success of the project relies on it : if response time is short, the system is attractive, and traffic rises, what requires even better performances.

Also, there is no use in having an efficient mechanism to retrieve data if these are poorly consistent as it was the case in the previous, manual system (see paragraph 3.1). Above all, data used to retrieve documents must be reliable, in order for users to get exhaustive answers which they can rely on. This problem can be solved by a DBMS, which imposes strongly structured data and strict controls on data entry.

A third problem is often neglected : data security, including access restriction and data protection at the storage site and during transfers. Manual systems can ensure reliable and complex security controls, while most automated systems do not give enough attention to these questions. It is essential to include robust security mechanisms in a system broadcasting thousands of (possibly classified) documents on a network.

2.4 General objectives of the project

There are two medium-term objectives in the SARDE project :

- In 1986-1987, the French Telecommunication Network is to be equipped with an operational system, which should be of interest for all operators and maintenance staff;
- the application should play a demonstrative role, showing that several techniques are realities, more particularly wide band networks. It is expected to be followed by several industrial developments using these new

technologies.

A prototype of the system will be first developed, demonstrated and evaluated. Then industrial companies will realize the final system. Only the prototype phase is scheduled today, and will be finished by the end of 1984.

3 Results of the preliminary analysis

Many problems arose during the preliminary analysis of the SARDE project, which could not be solved by commercially available DBMSs. We had to investigate several techniques from different fields. Even though the article focuses on data base related issues, each of the major fields of the SARDE project is briefly described.

3.1 Organization against inconsistency

The information contained in the current technical documentation database concerns three hundred thousand documents, representing several millions of pages. Consistency is far from perfect in such a large amount of data, which are produced and manipulated by so many people (typically, the evolution of a product is not recorded incrementally, or dates are wrong, or a lot of "mandatory" informations are omitted). A new organization is needed in order to avoid inconsistencies in the information system itself. The final system will need organization techniques, on a nationwide scale, combined with large database know-how.

The experience gained with the existing document delivery system shows that a powerful database server must be created. At this central site a reduced, specialized staff will be able to carry out database updates and to prevent inconsistencies in new data at the same time. However it will not be possible to eliminate errors or omissions present in the millions of pages belonging to the pre-existing documentation.

The present information system relies on the know-how of users who deal with the incoherences presented above. Ideally, the electronic system should minimize the number of transactions leading to some desired result. But since the existing documents cannot be cured, it is not always possible to get all the desired information with only one query. Thus it seems reasonable to provide users, who are experienced technicians, with several paths for locating the desired information, and with operations for combining results from different previous searches.

3.2 Information modelling

The information system contains different kinds of data :

- diagrams, drawings;
- computer programs;
- technical notes of product specifications;
- files listing references to other documents, for a given state in the evolution of the products which they describe;
- special documents recording changes in other documents.

Each kind of document provides specific, partly independent, semantics. Two other features make even more difficult the modelling process using conventional database models :

- equipments change often : manufacturers build new releases of products, providing either the same external functions after internal changes, or enhanced, different behaviour; these modifications are frequent for any given device, and every new version brings changes in all the documents referencing its specification note; it is essential that final users can have access to all, or to a selected part of, the documentation concerning the evolution of a certain product from one release to another; thus, documents are retrieved (among other paths) by the characteristics of the product they are related to, and/or by dates or version numbers.
- even though information must be structured, and retrieved, according to data items and semantic associations, the system must cope with a very heavy heritage from the past : some data are false, or absent; others do not follow the same coding rules, like reference codes and titles; as a matter of fact, most of the documents have been created by the services manufacturing the products, where reference and indexation rules have been interpreted in several different ways.

Whatever the nature of an information is, it must be in the database, and accessible for authorized users. Consequently :

- all the documentation (diagrams, computer programs, specification notes, update specifications, etc) is stored without interpretation, like images of the

original paper pages; these are delivered to the users as they have been stored; the acquisition process is presented in the next paragraph, the presentation model is described in section 4;
- semantics about documents are stored as structured data into a reference database; these are used to retrieve documents with search expressions on the nature, indexation, successive versions, etc, of the documents.

3.3 Image management

Initially, the documentation is available on paper or microforms. Three functions are necessary in order to manage the images of the original pages :

- digitalization
- high rate compression
- implementation of a complete acquisition system being able to operate at a reasonable cost on the voluminous documentation.

Digitalization is necessary in order to transfer the documentation from its present form to an electronic transmission compatible media. In the future, texts will be prepared with text processing or photocomposition tools, and technical documentation will be composed using CAD systems. For the existing documentation, scanners will be used, giving a 200 to 400 dots per inch resolution in order to provide images of good quality.

Compression of digitalized documents is mandatory. An 8"x11" page contains 4 Mbits, and the usual facsimile compression method is inadequate with regard to storage and transfer issues (with this method, an efficiency of 8 for unidimensional coding is typical, and up to 15 with bidimensional coding). A combination of raster to vector transformation, character recognition and residual areas facsimile compression will be applied to images, giving an improvement of 5 to 10 times of the usual compression rates. Although such a combination is not in common use, it is efficient when applied to big quantities of documents.

The acquisition process of documentation brings up the problem of its indexation. The whole technical documentation is contained in several hundred thousands of documents, each of them being identified with a cartouche, referencing the manufacturer, the system and the type of switching center. Also it contains the reference code of the document and its title. It is quite impossible to apply a manual indexing data entry for all the documents. An automatic pattern recognition subsystem must analyze the contents of the cartouche enclosed in every document. Its input is the digitalized image of the document, where the cartouche must be identified and analyzed. Its outputs are the characteristics listed above (reference code, etc). A control can be made by a team of documentalists 'a posteriori', with an application showing the cartouches and the corresponding, automatically built, index entries.

Even after high compression, 5 millions of pages necessitate about 50 Gigabytes of storage capacity, that is to say 50 optical disks. The evolution rate of the documentation is constant, and no deletion is possible since it is an archival system. This last characteristic indicates that writable only once optical disks (WOOD) are an affordable solution, compared to magnetic disks (from 10 to 100 times cheaper). In order to use them with a short response time, the architecture of the WOOD management subsystem must be designed to prevent bottlenecks in communications and to provide direct access to any given image.

3.4 User interface

The users should be able to get any required information as if it were near hand at their own maintenance site. The distinction between images and semantic bibliographical data is not visible for them. They query the reference database with a dedicated interface, and images are displayed on the screen of their workstation. These aspects (external multi-type database model, query language and hardware environment) are developed in section 4.

3.5 High speed communication

A very fast electronic transport service is necessary to bring digitalized images in a few seconds from the central site to the user's workstation. Over the 2000 maintenance centers, only a few hundred centers will probably be connected at any given time to the central system, 50 are estimated to be in an active dialogue phase, and there should be more than ten simultaneous transactions being processed by the server. About three images should be sent at the same time by the server, interleaved with queries to, and answers from, the references database. A 64 Kbps network will be adequate for these requirements (this is a clear demonstration that high speed computer networks are needed without delay for a near third generation of DBMSs).

3.6 Structured data management

In order to model the semantics of the reference database at a conceptual level, several notions had to be added to the relational model (CODD71). Distinct entities may convey the same semantics, for instance some diagrams, specification notes, technical documents, etc. The notion of generalization (SMITH78) helps in specifying only once the characteristics of a root leading to several specific entities which share these same characteristics. Also, there are n-ary associations which identify a group of entities as a whole, like users or groups of users authorized to access a document. The notion of aggregation minimizes the effort of describing such associations, and ensures consistent schema updates.

The evolution of the documentation has been taken into account with a list of editions for every document, and a list of pages for every edition. Since only some pages are changed from one edition to another for a given document, it is possible to store on WOODs only one instance of every page, whatever edition(s) it belongs to. Eventually, the omissions or errors, unavoidable in the original data, are taken into account by the system, thanks to the tolerance of the database schema.

The conceptual model was mapped on a relational model, as a provision for executing unscheduled queries. But the volume of information is not compatible with the way commercially available relational DBMSs handle data on a central site. Conventional systems are unable to process ten queries at a time on a five hundred Mbytes relational database, with a response time between three and ten seconds, according to the complexity of the query. Our final users will experiment an entirely computerized system for retrieving technical documentation, and a short response time is mandatory. It is not possible to build little databases for equipments of each site, because, at any given location, almost the whole documentation is required. So, only a large "centralized" database is possible. In these conditions, special storage devices must be used at the central site, which are able to accelerate physical access to data, and to process quickly some costly functions like relational operators. Cache memories, electronic disks, intelligent programmable disks and database machines must be combined for the most efficient throughput. One solution using COPERNIQUE's devices is underway for each phase of the project.

3.7 Security

The database contains information on patented products, and undesired access must be prevented. For every query, the documents indicated and shown to the user must be strictly among those which he is authorized to see. Also, the use of a network brings the problem of protection during data transfers. These requirements necessitate enhanced protection mechanisms compared to what is commonly proposed by commercial DBMSs. The solution is emphasized in section 7.

4 General architecture of the prototype

Usually, an application database software is made of different components, which are functionally well separated, but in fact located in the same host computer. In our case, three independent machines are involved in the query process. The workstation supports the query acquisition and the answer presentation. The CII-HB Level 6 host translates the external query in an internal form, and participates to some phases of its execution. The DIRAM 32 executes powerful searches and sends full answers to the host.

4.1 Workstation

The workstation is a multiprocessor microcomputer, partly running under a Unix-like system. It is able to support a complete acquisition and validation of the query, using screen forms. A formatted and validated query is sent to the host, which answers with a result file. This file can contain a simple view of a result screen, but generally it contains more information. The local software displays its contents, according to the user's presentation requirements and to screen size constraints. This procedure and the local treatments on retrieved data save many exchanges with the host and simplifies its software.

4.2 External requests translation

The external form of the query (e. g. "access_document <#ref> -list_of_editions -list_of_pages") must be translated in internal requests. This could be done in two phases. The query would be first translated in a conceptual form, which expresses the algorithm, with no care about access paths implementation. Then the conceptual to internal translation would be executed. However the purpose of the first prototype is a feasability study, and the realization of a complete, automatic external-conceptual mapping is not mandatory. Consequently, external requests are actually implemented with programs using directly the internal level commands. The software developed for this application is equivalent to a dedicated DBMS, but any internal change leads to necessary changes in the request programs. In order to minimize maintenance efforts, the internal level of service provides a request specification language that helps in changing application programs (see section 5).

4.3 Internal level of service

In the first phase of the project, physical access to data is accelerated with an intelligent, programmable disk, the DIRAM 32. The internal level of service is made of a set of commands working on the relational structure resulting from the mapping of the conceptual model of the reference database.

The dealine to first experiment the system is December 83, and time-dependent choices have been made concerning the commands to be programmed. It has been decided to use existing funtions (like Level 6 FMS) as much as possible, and to program accelerators to relational implementation. The central server is built around a CII-HB Level 6 connected to a DIRAM 32. GCOS MOD 400 FMS functions as index management and relative access files are used to implement entities, associations between entities and access by key values. DIRAM 32 is programmed to do selects, projects and set-oriented navigation at the same time (see section 5).

In the second phase of the project, all possible requests will be addressed by a DORSAL 32 database machine providing at least the same transaction throughput, but managing sharing and resiliency too (see section 6).

4.4 Updating the database

The image database is an archival system updated by additions of new images on WOODS. The reference database processed by the internal commands (GCOS and DIRAM functions) is built from a twin database already managed on the Level 6 by SOCRATE, the French CODASYL-compliant DBMS developed by Abrial (ABRI74). The internal structure of this twin database cannot be efficiently exploited by any internal command at the level of the DIRAM. So it is necessary to load the reference, relational database from the SOCRATE database. This choice minimizes the development costs : the internal structure of the relational database makes this operation easy to program and to run, while SOCRATE includes all the mechanisms ensuring data entry controls and resiliency against incidents.

5 The DIRAM 32 programmable, intelligent disk

5.1 Product overwiew

The main bottleneck for computational power is I/O delays. In order to solve this problem, storage systems have been developed which include cache memory and local intelligence. The DIRAM 32 belongs to this recent generation of intelligent disks. Its basic multiprocessor architecture contains :

- one main processor managing the storage unit and the communication at the logical level between the host computer and the other DIRAM 32 components;
- one disk formatter managing from 1 to 16 disk units;
- a cache memory (up to 16 Mbytes) implemented with passive modules;
- a pair of interconnected couplers, one in the DIRAM and the other one in the host computer.

This configuration is typical for I/O bound applications which need to be boosted. The cache memory allows to choose operation parameters like anticipated block reads, LRU replacement strategy vs locking of certain pages of the working set in the cache, postponed writes following the lift strategy, etc.

Standard mechanisms ensure data integrity. Postponed sector writes are stored in RAM powered by a 30-minute battery in case of incident in the DIRAM 32. Each disk can be physically doubled by another, and still be considered as only one logic unit. In this last configuration, read actions are processed on the disk which has the best angular position, and each write action is done on both physical units.

This set of facilities is transparent to the host computer. The DIRAM 32 is physically and logically plug-to-plug compatible with several computers (Honeywell Level 6, Gould SEL 32, etc). It replaces the standard disk coupler in these systems, and re-uses the existing disk units, as long as they provide a SMD interface.

For special applications processing very large amounts of data, one or more of an additional kind of component can be added to the DIRAM 32. This component is called "active module", and comprises one Motorola 68000 microprocessor with a 256 Kb local memory, where application dependent functions can be loaded and run.

An active module has direct read access to disk units managed by the formatter. If the host computer communicates a disk area to an active module, this one can process a dedicated treatment on sectors from the designated area, then send the results to the host via the I/O channel. Up to seven active modules can work in parallel in a DIRAM 32. Active modules have already been used in many applications, from image generation (real-time animation for pilot training) to electronic telephone directory accessed on a regional scale by thousands of consumers from home Videotex terminals.

5.2 Programming DIRAM 32 for the SARDE project

The reference database is implemented using the relational model. Entities of the conceptual level are mapped on relations, n:p associations also. Generalizations are implemented through attributes containing the nature of the entity implemented by a given tuple. Semantic properties of aggregations are extended to all aggregated entities.

In order to enhance set processing, additional attributes have been included in relations representing associations : they provide a structure of list for all tuples of a relation tied to one tuple of (possibly) another relation. Thus the structure of the database is still entirely relational, and information relative to links is separated from original data relations. But thanks to this dichotomy, the database can be processed either by navigation through the logical links with predefined transactions, or by relational set operators for occasional queries.

The active modules of the DIRAM 32 used in the prototype of the SARDE project run internal commands which mainly consist of table- and data-driven automata :

- a) from one given tuple in a relation,
- b) process all the tuples of a (possibly different) relation participating to a link tied to this root, and apply complex selection criteria and multi-attribute projections,
- c) possibly repeat recursively the process for each selected tuple (go to a) or terminate.

The level of service of the DIRAM 32 functions has been formalized with an internal command language. This language allows to phrase the internal requests before coding the application programs which will activate DIRAM 32 active modules. Also several options are offered : some of them enhance programmer productivity (like conditional processing steps or self-test mode), others improve response time (like attribute duplications in the relations implementing semantic links).

5.3 Activating commands in DIRAM 32

Several commands can be sent through the I/O channel by the Level 6 application to the DIRAM 32, where they are interpreted by the main processor, then executed by the active modules :

- LOAD an active module with the code of a given function;
- START the program with input parameters;
- RESTART or CONTINUE the program if it has been temporarily stopped (several functions can run in one active module as coroutines of host programs);
- HALT terminates the execution of the function.

5.4 Performance evaluation

A more sophisticated physical organization could have been designed for the reference database. But the programming effort for loading, updating and searching such a complex organization would not have been compatible with experimentation deadlines. Performance evaluations show that the active module throughput is enough to deal with several simultaneous queries. The limits of the first configuration (one active module, 2 Mb cache memory) can be enhanced if more active modules and cache memory are added.

6 A complete database machine : the DORSAL 32

6.1 Product overview

As its name indicates, the DORSAL 32 is a back-end computer entirely dedicated to the management of databases (ARMI81). This commercial product is connected to host computers which run the applications. It provides an interface conforming itself to the recommandations of the CODASYL (CODA71), and a SQL-like relational interface (ASTR75). Both are mapped on a unique internal data model, and the same base can be accessed and updated concurrently using both modes.

The architecture of the DORSAL comprises two main machines, each of them being multiprocessor : the logical management unit (LMU) runs the DBMS code, and the physical management unit (PMU) is in charge of the disks (see figure 6.1).

The LMU contains :

- one main processor,
- from 1 to 7 coprocessors (Motorola 68000),
- one back-end control unit (BCU, Motorola 68000), from which maintenance and statistics are possible,
- a 256 Mb address space of MOS memory;
- a pair of interconnected couplers, one in the LMU and the other one in the host computer.

Each processor has access to the segmented, paged memory through a MMU verifying capabilities, and providing a hierarchical topology.

The PMU is a DIRAM 32 with several active modules used for relational operations.

Different configurations are possible :

- the disks can be connected to two DORSAL 32, the active one and the back-end machine;
- one DORSAL 32 can be connected to 4 host computers;
- disks can be physically doubled for higher resiliency (see description of the DIRAM 32 standard facilities).

From a software point of view, an application program is written using either the network or the relational interface. It works on a subschema of the database. It is precompiled (in COBOL, FORTRAN), compiled and run. Each transaction running in the host computer begins a dialogue with the LMU of the DORSAL 32. The LMU manages communications between the DORSAL 32 and the host (presentation and session layers) and provides utility programs; it creates one run-unit in a coprocessor per transaction opened in the host and takes care of concurrent accesses, journalization and recovery. Coprocessors run access methods (the DBMS kernel) for the transactions which have been assigned to them.

The PMU is used by the DBMS kernel as a disk controller with accelerating cache memory. The active modules run projects, selects and joins on criteria which do not need to correspond to predefined paths in the database schema.

Many interesting features of the DORSAL 32 cannot be completely described here, as for instance data modelling, locking policy, recovery mechanisms, set-oriented navigational functions, transaction processing, query optimization, data flow scheduling, subschema static and dynamic definitions, data dictionnaries or dynamic schema modifications. These functions reflect recent advances in database technology, and are proposed in the DORSAL 32 commercial product.

6.2 Programming the application with DORSAL 32

The second phase of the SARDE project will lead to a real scale implementation. It will need an integrated environment for updating and searching the reference database. This can be done with DORSAL 32 including several coprocessors and active modules.

Many advantages can be found in such a solution : the conceptual database semantics are taken into account by a strongly structured schema, independently of their physical implementation. Resiliciency mechanisms are automatic. A complete application development methodology is available with subschemas, precompilers, etc.

7 Security

The national telephone switching network is the most widely used communication media in France. With the SARDE project, all the documentation related to it will be gathered in a central server, which will become de facto a critical point for the security of the country. Since this documentation can be accessed by any distant workstation, protection must be present at the central site, at every maintenance center and on the data links. Physical measures can be combined with organizational decisions and algorithmic protections.

7.1 Distributed cryptosystem

We can consider the whole security system as a cryptosystem, including information flow specifications, coding techniques and their implementation. Four problems are inherent to conventional cryptosystems :

- it is almost impossible to be sure that

non-numerical encipherment techniques (based on susbstitutions and transpositions) cannot be broken;
- the delivery of (de)coding keys to the users must be itself protected (... with a cryptosystem);
- these keys must be kept secretly;
- when a user communicates, there should be a way to authentify its identity, and to prevent unauthorized intrusions.

The first problem is due to the relative weakness of usual coding techniques compared to the computational power of decipherment centers. The three others are consequences of hasardeous organizations for communicating secret keys. All four problems have received a global solution with public-key cryptosystems (DIFF76). These are based on a separation between the coding and the decoding functions, such that knowing one, it is not possible with present computers to find the other one (to break the code is a NP complete problem). A pair of such functions is attached to each user, and he makes one of them public. Every communication between members of the cryptosystem permits to protect data and to authentify the correspondent's signature (RIVE78). Encipherment functions are easy to compute, and the corresponding overhead can be reduced if functions are changed regularly. This solution can be implemented at low cost on the different computers involved in data transmission. It should be applied for any transfer of classified information, but perhaps not for images which have been compressed by several, combined techniques. The same kind of encipherment function could also be used to protect digital images, but real-time coding and decoding processes on 100 Kbits images would necessitate special hardware devices.

7.2 Document typology and user's classification

Every user must sign on when he logs on the system, with a password that only he can change. After this identification phase, he will have access to informations according to his general access rights.

Documents and keywords are classified. Auhorizations are delivered to users and groups of users for accessing documents which they are entitled to know of. The application filters the information (documents, editions and keywords) according to the user authorization, taking into account its group membership. This control is implemented as a query modification in the relational model (STONE78) and executed by the DIRAM 32 as any elementary select.

8 Conclusion

The SARDE project conducted by CNET, the French National Center for Research in Telecommunications, leads to the implementation of an entirely electronic documents delivery system concerning the technical documentation on the national telephone switching network. The enormous volume and the complexity of this documentation necessitate a combination of several techniques from different fields : organization and systems analysis, information modelling and presentation, image management, high level graphical user interface, high speed communication links, efficient structured data management and security in a distributed system.

A first phase, scheduled for 1984, will allow to evaluate the feasibility of managing the totality of document references with a relational database accessed through an intelligent disk, the DIRAM 32. Later, a second phase will bring the whole application to a complete database machine, the DORSAL 32. Besides the efficiency of query processing, a special emphasis will be given to external interface after feedback from final users, and also to security issues as soon as the preliminary experiments will be completed on the database.

ACKNOWLEDGMENTS

The authors want to thank the referees who provided several remarks that helped in clarifying the presentation of some technical aspects of the SARDE project; they are also particularly indebted to Elaine Penny for her encouragements and comments while writing this article, and to Christian Soulignac for text editing.

REFERENCES

ABRI74
J.R ABRIAL, Data Semantics, in Data Management Systems, Kimbie & Koffeman (eds), North Holland 1974.

ANSI77
D. TSICHRITZIS & A. KLUG (eds), The ANSI/X3/SPARC DBMS Framework, Tech Note 12, Computer Systems Research Groups, University of Toronto, July 1977.

ARMI81
J.P ARMISEN, J.Y CALECA, A commercial back-end system, Proc. VLDB 81.

ASTR75
M.M ASTRAHAN, D.D CHAMBERLIN, Implementation of a Structured English Query Language, CACM Vol 18, No 10, October 1975.

CODA71
CODASYL COMMITTEE ON DATA SYSTEM LANGUAGES, CODASYL DBTG Report, ACM, New York, 1971.

CODD70
E.F CODD, A Relational Model of Data for Large Shared Data Banks, CACM, Vol 13, No 6, June 1970.

DIFF76
W. DIFFIE, N. HELLMAN, New Directions in Cryptography, IEEE Transactions on Information Theory, Vol 6, November 1976.

RIVE78
R. RIVEST, A. SHAMIR, L. ADLEMAN, A method for obtaining digital signatures and public-key cryptosystems, CACM, Vol 21, No 2, February 1978.

SMIT77
J.M SMITH, D.C.P SMITH, Database Abstractions : Aggregation and Generalization, ACM TODS, Vol 2, No 2, June 1977.

STON75
M. STONEBRACKER, Implementation of integrity constraints and views by query modification, Proc ACM SIGMOD Conf San Jose CA, May 1975.

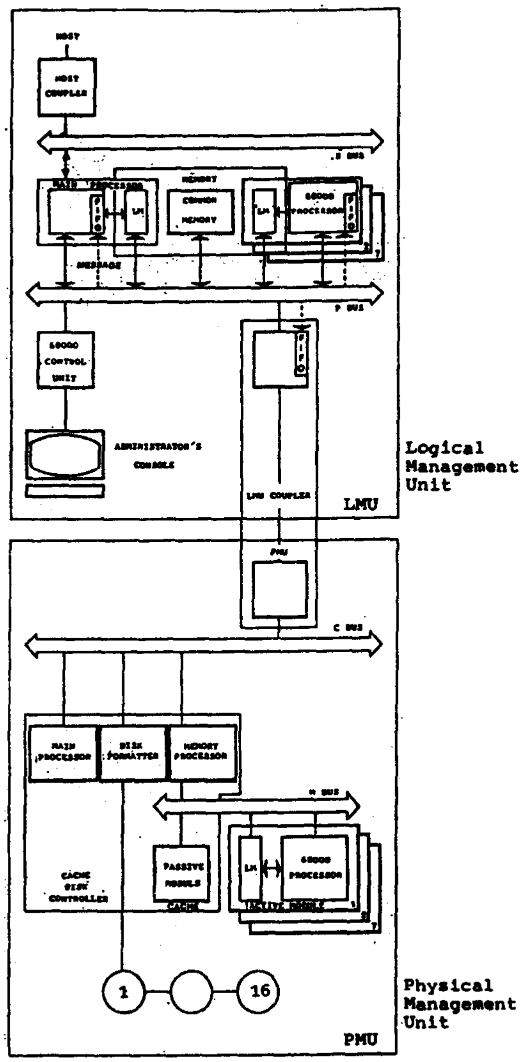

Figure 6 - 1 : Architecture of the DORSAL 32.