

*Project-Team ATLAS**Complex Data Management in Distributed
Systems**Rennes*

THEME 3A

The logo consists of a large, stylized, light blue 'A' on the left and a large, stylized, light blue 'R' on the right. The word 'Activity' is written in a light blue serif font, positioned between the two letters. The word 'Report' is written in a light blue serif font, positioned below the 'R'. A horizontal line is drawn across the middle of the 'A' and 'R'.

2003

Table of contents

1. Team	1
2. Overall Objectives	1
3. Scientific Foundations	2
4. Application Domains	2
5. Software	3
5.1. ATL (Atlas Transformation Language)	3
5.2. $Find_{AGE}^m$	3
5.3. RepDB*	3
5.4. SaintEtiQ	4
6. New Results	4
6.1. Database summaries	4
6.1.1. Summaries for Image Clustering	4
6.1.2. Summarization Technique for TV Recommendation Systems	5
6.2. Model management	5
6.2.1. ATL (Atlas Transformation Language)	5
6.3. Multimedia data management	6
6.3.1. Personal image collection management from mobile devices	6
6.4. Distributed data management	6
6.4.1. Data replication in cluster systems	7
6.4.2. Distributed data processing	7
7. Contracts and Grants with Industry	8
7.1. Leg@net RNTL Project	8
7.2. Microsoft Research Contract	8
7.3. Domus Videum RNTL Project	8
8. Other Grants and Activities	8
8.1. Regional Actions	8
8.2. National Actions	8
8.3. International actions	9
9. Dissemination	9
9.1. Animation of the scientific community	9
9.2. Editorial Program committees	9
9.3. Invited Talks	10
9.4. Teaching	10
10. Bibliography	10

1. Team

Head of Project Team

Patrick Valduriez [Director of Research, INRIA]

Vice-head of Project Team

Noureddine Mouaddib [Professor, University of Nantes]

Faculty Members, University of Nantes

Jean Bézivin [Professor]

Marc Gelgon [Associate Professor]

José Martinez [Associate Professor (HDR)]

Esther Pacitti [Associate Professor]

Guillaume Raschia [Associate Professor]

Technical Staff

Gaëtan Gaumer [Engineer, CDD]

Administrative Assistant

Marie Pivaut [Secretary, CDD]

Ph.D. Students

Reza Akbarinia [Fellowship Iranian gov., since oct. 2003]

Alexandre de Assis Bento Lima [Fellowship CAPES, Brazil, since oct. 2003]

Cédric Coulon [Fellowship LIP6, RNTL contract]

Frédéric Jouault [Fellowship CIFRE, TNI, since oct. 2003]

Erwan Loisant [Fellowship Monbusho, Japan]

Vidal Martins [Univ. PUCPR, Brazil, since oct. 2003]

Antoine Pigeau [Fellowship MENRT]

Régis Saint-Paul [Fellowship MENRT]

Andre Victor [Fellowship CNPQ, Brazil, since oct. 2003]

Amnel Voglozin [Fellowship MENRT]

2. Overall Objectives

Today's hard problems in data management go well beyond the traditional context of Database Management Systems (DBMS). These problems stem from significant evolutions of data, systems and applications. First, data have become much richer and more complex in formats (e.g., multimedia objects), structures (e.g., semi-structured documents), content (e.g., incomplete or imprecise data), size (e.g., very large volumes), and associated semantics (e.g., metadata, code). The management of such data makes it hard to develop data-intensive applications and creates hard performance problems. Secondly, data management systems need to scale up to support large-distributed systems (such as the Internet or cluster systems) and deal with both fixed and mobile clients. In a highly distributed context, data sources are typically in high number, autonomous and heterogeneous, thereby making data integration difficult. Third, this combined evolution of data and systems gives rise to new, typically complex, applications with ubiquitous, on-line data access: virtual libraries, virtual stores, global catalogs, services for personal content management, services for mobile data management, etc.

The general problem can be summarized as complex data management in distributed systems. The Atlas group addresses this problem with the objective of designing and validating new solutions with significant advantages in functionality and performance. To tackle this objective, we separate the problem along four main dimensions. The theme "database summaries" addresses the issues of data abstraction from large size databases. The theme "model management" addresses the issues of data abstraction from complexity. The theme "multimedia data management" deals with efficient and personalized access to multimedia data. Finally, the theme "distributed data management" addresses the problems of data replication and distributed query processing with complex data.

3. Scientific Foundations

Key words: *Data management, database, distributed database, multimedia, summaries, fuzzy logic, meta-modelling, distributed systems.*

The Atlas group relies on solid, complementary foundations in data management, meta-modelling and distributed systems. Data management is concerned with the storage, organization, retrieval and manipulation of data of all kinds, from small and simple to very large and complex. It has become a major domain of computer science, with a large international research community and a strong industry. Continuous technology transfer from research to industry has led to the development of powerful DBMSs, now at the heart of any information system, and of advanced data management capabilities in many kinds of software products (application servers, document systems, directories, etc.).

The fundamental principle behind data management is data abstraction, which enables applications and users to deal with the data at a high conceptual level while ignoring implementation details. This principle has been achieved by providing advanced capabilities such as high-level query languages, automatic query processing and optimization, data and meta-data modelling, data structures for supporting complex data, transactions, etc.

An important kind of data that we address is multimedia data, such as image, audio or video, which is quite different from structured data and even semi-structured data, such as textual documents. Multimedia data is typically translated into sets of features which are mathematical properties obtained from signal processing and analysis techniques, e.g., colour histograms for images. If we can rely on such a theoretical background for the first stage of the translation process, there are still open problems in adapting these low-level descriptions into something more manageable by data management systems.

Database query languages are typically based on first order logic. To allow for more flexible manipulation of large quantities of data, we rest on fuzzy logic to summarize data.

Meta-modelling and meta-data have been studied by the database community for a long time. It is interesting to witness the impact of similar principles in software engineering. So-called meta-models are used today to define domain specific languages that may help capturing the various “aspects” of software-intensive systems. Models are no more viewed as “contemplative” artefacts, used only for documentation or for programmer inspiration. In the new vision, models become computer-understandable and may be applied a number of precise operations. Among these operations, model transformation is of high practical importance to map business expression onto executable distributed platforms but also of high theoretical interest because it allows establishing precise correspondences between various representation systems without ambiguity. For these reasons, model engineering has become the subject of much collaboration between the data management and the software engineering areas.

Finally, the Atlas group considers data management in the context of distributed systems, with the objective of making distribution transparent to the users and applications. Thus we capitalize on the principles of distributed systems, in particular, large-scale distributed systems such as clusters, grid, and peer-to-peer (P2P) systems, to address issues in data replication and high availability, transaction load balancing, and query processing.

4. Application Domains

Key words: *Application Service Provider (ASP), large decision-support application, multimedia personal database.*

Complex data management in distributed systems is quite generic and can apply to virtually any kind of data. Thus, we are potentially interested in many applications which help us demonstrate and validate our results in real-world settings. However, data management is a very mature field and there are well-established application scenarios, e.g., the On Line Transaction Processing (OLTP) and On Line Analytical Processing

(OLAP) benchmarks from the Transaction Processing Council (TPC). We often use these benchmarks for experimentation as they are easy to deploy in our prototypes and foster comparison with competing projects.

However, there is no complete benchmark that can capture all the requirements of complex data management. Therefore, we also invest time in real-life applications when they exhibit specific requirements that bring new research problems. Examples of such applications are Application Service Provider (ASP), large decision-support applications or multimedia personal databases.

In the ASP model, customers' applications and databases (including data and DBMS) are hosted at a provider site and need to be available, typically through the Internet, as efficiently as if they were local to the customer site. Thus, the challenge for a provider is to manage applications and databases with a good cost/performance ratio. In Atlas, we address this problem using a cluster system and exploiting data replication and load balancing techniques.

Large decision-support applications need to manipulate information from very large databases in a synthetic fashion. A widely used technique is to define various data aggregators and use them in a spreadsheet-like application. However, this technique requires the user to make strong assumptions on which aggregators are significant. In Atlas, we propose a new solution whereby the user can build a general summary of the database that allows more flexible data manipulation.

A major application of multimedia data management that we are dealing with in Atlas is multimedia personal databases which can help retrieve and classify personal images stored on a hard disk or CD-ROM. Such domestic applications, extended to the video medium mainly, could be integrated into "intelligent" TV sets. Currently, the integration of multimedia is effective only with images. From the usability point of view, open issues are the effective combination of various medias and the adaptability of the indexing process to a specific task or application domain.

5. Software

5.1. ATL (Atlas Transformation Language)

Participants: Jean Bézivin, Frédéric Jouault, Patrick Valduriez.

ATL is a transformation-based model management framework, with metadata management as the main application. The ATL language is designed to be general and abstract. We plan to use it to compile transformations to many different target languages including XSLT and XQuery. The ATL design strives to be consistent with the MDA standards, in particular MOF/QVT. The ATL system is being implemented in Java, and we plan to port major transformation components to the .Net platform. We plan to release the model management system and its components for metadata management as Open Source Software under a non restrictive license such as BSD. Furthermore we are planning to make our components publicly available to the ObjectWeb consortium.

5.2. $Find_{AGE}^{Im}$

Participants: José Martinez, Erwan Loisant, Régis Saint-Paul.

$Find_{AGE}^{Im}$ is an image search-by-content system. Currently, it is fairly complete from the architectural point of view. It provides three different ways to query images by content: formal querying, interactive querying and browsing. Formal querying is based on the traditional querying approach developed for structured DBMSs. The interactive querying process of $Find_{AGE}^{Im}$ is based on the information retrieval querying process, i.e., when manipulating noisy data it is hardly possible to write down immediately the correct query, if at all. Browsing is a more efficient and effective way to retrieve rapidly visual information such as images. Currently, a first complete implementation variant of this architecture is on the way to be achieved.

5.3. RepDB*

Participants: Cédric Coulon, Esther Pacitti, Patrick Valduriez.

RepDB* is an Open Source Software data management component for replicating autonomous databases or data sources in a cluster system. It has been developed in the context of the Leg@net RNTL project. RepDB* supports preventive advanced data replication capabilities which are independent of the underlying DBMS. It uses general, non intrusive techniques. It is implemented in Java on Linux and supports various DBMS: Oracle™, PostGresQL and BerkeleyDB. It has been validated on the Atlas 8-node cluster and another 32-node cluster at INRIA-Rennes.

5.4. SaintEtiQ

Participants: Noureddine Mouaddib, Gaëtan Gaumer, Guillaume Raschia, Régis Saint-Paul, Amenel Voglozin.

SAINTETIQ is a data summarization system providing synthetic user-friendly views over large databases. The fuzzy-set based representation of summaries provides an effective way of dealing with uncertainty in data, and natively supports flexible queries. A user-centric approach of a summary-oriented knowledge discovery process has been integrated into the prototype. We also enhanced the implementation with a set of tools to generate the background knowledge required for the summarization process. Finally, a complete graphical user interface has been developed to support the user manipulating and browsing data, background knowledges and summaries.

6. New Results

6.1. Database summaries

DBMS has become a very mature technology that is ubiquitous in information systems. Over time, the extensive use of DBMS technology has had major consequences in large organizations: the production of very large databases, the production of heterogeneous databases, and the increasing requirement of diverse applications to access those very large, heterogeneous databases. This creates difficult technical problems which get worse as DBMS technology improves and is more able to produce very large, heterogeneous databases. The SAINTETIQ system provides a novel solution for representing, querying and accessing large databases. Recent work focused on the exploration of concrete use-cases in two distinct application areas. First, we developed an approach relying on SAINTETIQ summaries to support an efficient navigation process of image databases. Secondly, we addressed the problem of User Profiling through a real-world application concerned with Personal Video Recorders and TV program filtering.

6.1.1. Summaries for Image Clustering

Participants: Erwan Loisant, José Martinez, Noureddine Mouaddib, Guillaume Raschia, Régis Saint-Paul.

Querying image databases with similarity searches and relevance feedback has been largely investigated in the literature. In contrast, browsing has not been much studied. We proposed a browsing technique based on the database summaries produced by SAINTETIQ and a Galois' lattice used by $Find_{AGE}^m$. This technique helps avoiding the high cost incurred by Galois' lattices alone. The result is a kind of hypertext of images that combines classification and visualization in a high-dimensionality space.

Visual browsing allows the user to navigate through the entire image. This approach is very intuitive and no complex manipulation of image features is required. Also, images can be compared and classified much more precisely because expensive calculations are done off-line. We proposed to divide the navigation process in two levels: general and specific.

The general level is a navigation between collections of images. It is based on a Galois lattice to organize homogenous sets of images and offers an intuitive way to navigate between them. This data structure offers several advantages : it achieves a kind of multidimensional indexing; the speed at which the image database can be browsed is independent of the number of images; and this structure can be used for physical indexing of large image databases.

The specific level is a navigation inside a particular set of images to further discriminate their individual characteristics with a hierarchical presentation. The formation process of the homogeneous sets of images and their descriptions relies on the SAINTETIQ process. From the large original database of images and their features, SAINTETIQ builds a hierarchy of summaries. A selection process is applied to select a fixed number of the most interesting summaries which will be used to produce the Galois' lattice. The selection process also ensures that the selected summaries cover the entire database

6.1.2. Summarization Technique for TV Recommendation Systems

Participants: Marc Gelgon, Nouredine Mouaddib, Guillaume Raschia.

The increasing number of satellite and cable television channels is resulting in a soaring number of broadcast programs available to viewers. Thus, there is a need to develop schemes dedicated to home multimedia platforms, that assist users in the selection of programs of interest, i.e., that attempt to make relevance to user-based filtering of the program stream. In order to characterize the program content, one may either extract metadata from the audiovisual material itself, or rely on accompanying metadata. We opted for the latter, given that such descriptors are progressively being made available in practice, either in the program stream, e.g., Digital Video Broadcast Service Information (DVB-SI), or as program guides on the internet.

Our approach focuses on learning the user profile through the behaviour (viewing history) of the user. We in fact advocate a system that relies both on explicit, i.e., manually defined, and implicit profiling. The present work proposes a novel technique for tackling automatic learning of the user profile and for program filtering. This is carried out through the generation of a tree-based organization of SAINTETIQ-like fuzzy concepts. The learning task is performed on-line, summarizing the metadata received by the system and characterizing the user interests in this tree. This offers a great flexibility in the design of profiles. Another feature of the proposed scheme is that it is well suited to handling, in the same manner, both numerical and symbolic metadata. Indeed, both forms are likely to appear in programs.

6.2. Model management

A model is a structure that represents a design artifact such as a database schema, an interface definition, an XML type definition, a UML model or a Web document. Developers of information systems must typically deal with different models and perform transformations between models. Examples of transformations are: mapping heterogeneous data source descriptions in a global schema to perform data warehousing, converting XML documents into HTML, or generating EJB or .Net component definitions from a UML model. Today, most of these transformations are still programmed using specific languages like SQL, XSLT or even Java or C. As information systems become more complex and need to support cooperation of heterogeneous applications and components, such manual development of models and transformations is no longer viable.

Model management aims at solving this problem by providing techniques and tools for dealing with models and model transformations in more automated ways. It has been studied independently for years by several research communities such as databases, document management, and software engineering. One of the major problems is the multiplicity of input and output format and transformations systems, e.g., from Latex to HTML or from UML to Java. There is much to gain if we could handle these various transformations with a coordinated family of languages. Recently, the need for more generic model management has been recognized. To contribute to this evolution, we are designing ATL (Atlas Transformation Language).

6.2.1. ATL (Atlas Transformation Language)

Participants: Jean Bézivin, Frédéric Jouault, Patrick Valduriez.

ATL is a transformation-based model management framework. One objective is to support reusability through composition and specialization. Although it is still too early to state precisely how this will be achieved, the systematic use of meta-models to define abstract syntax and semantics for models is a good basis. Meta-models that describe domain-specific languages may be made composable and comparable through the use of a common meta-meta-model (a set of constraints on the kind of directed graph we use together with a set of

utilities to perform general operations). There are several families of such meta-meta-models, one of the most prevalent being the OMG Meta-Object Facility (MOF).

As evidenced by the wide adoption of UML, standard languages are important for model management. Thus, ATL will capitalize on the work on Query/Views/Transformation (QVT), the transformation language being defined by the OMG as well as on recent work on generic model management platforms. The key design decision in ATL is that models and model transformations are treated as first-class citizens. In particular, ATL supports higher-order transformations which are proving to be very useful for transformation reuse.

We have defined the requirements for the ATL framework from which we have derived an initial ATL definition. Unlike XSLT or XQuery which deals with trees, ATL deals with graphs. We have defined several benchmarks to study the properties of ATL. For example a semantic-preserving transformation for translating XSLT into XQuery has been written in ATL. This is a transformation composed of a number of simpler ones. The meta-models defined to handle input or output are extensions of the basic XML meta-model. We plan to use this technique to interface with other technological spaces as well, e.g., with abstract syntax languages where EBNF import/export facilities will be similarly organized.

We cooperate with the Triskell group at IRISA which is also developing a transformation language called MTL. ATL follows a more declarative approach than MTL. In a future version, the ATL engine will use a pivot system called basicMTL to interface with model repositories.

6.3. Multimedia data management

The ability to store multimedia information in digital form has spurred both demand and offer of new electronic appliances (e.g., DVD players, digital cameras, mobile phones connected to the Web, etc.) and new applications (e.g., interactive video, digital photo album, electronic postcard, distance learning, etc.). The increasing production of digital multimedia data magnifies the traditional problems of multimedia data management and creates new problems such as content personalization and access from mobile devices. The major issues are in the areas of multimedia data modelling, physical storage and indexing as well as query processing with multimedia data. We pursued one research action: personal image collection management from mobile devices.

6.3.1. Personal image collection management from mobile devices

Participants: Marc Gelgon, Antoine Pigeau.

Extension of image retrieval systems to address personal image collections appears among emerging needs. In particular, mobile devices such as camera-equipped phones are an interesting case for content creation and retrieval. In this context, we have proposed an unsupervised technique for organizing an image collection, based on time and geolocation meta-data. The objective is to recover the natural spatial and temporal structure present in such a data set. This meta-data is indeed both reliable and appealing for spatio-temporal browsing, as shown by user needs studies. The proposal is formulated as statistical mixture model-based classification. Dedicated optimality criterion and expectation-maximization search technique have been devised. An incremental version of the scheme has been defined, as content creation and retrieval are not consecutive in this application. Since the data sometimes exhibits spatial or temporal clusters at multiple scales, we are currently defining an extension of the scheme to recover such a structure, with a view to hierarchical browsing. As in the results described in section 6.1.2, the structure obtained should help efficient management of the data.

6.4. Distributed data management

In a large scale distributed system, data sources are typically in high numbers, autonomous (under strict local control) and very heterogeneous in size and complexity. Data management in this context offers new research opportunities since traditional distributed database techniques need to scale up while supporting data autonomy, heterogeneity, and dynamicity. There are different distributed system contexts where we can study these problems, in particular Internet and clusters of PC. However, to yield general results, we strive to develop

common algorithmic solutions with the right level of abstraction from the context. Thus, we assume a peer-to-peer (P2P) distributed system architecture which is able to scale up to very large configurations. Given such an architecture, data consistency and the performance of data access are crucial. To address these general problems, we have pursued two complementary research actions on data replication in cluster systems and distributed data processing.

6.4.1. Data replication in cluster systems

Participants: Cédric Coulon, Esther Pacitti, Patrick Valduriez.

Clusters of PC servers provide a cost-effective alternative to tightly-coupled multiprocessors. They have been used successfully by, for example, Web search engines using high-volume server farms (e.g., Google™). However, search engines are typically read-intensive which makes it easier to exploit parallelism. Cluster systems can also make new businesses such as Application Service Providers (ASP) economically viable. To improve performance, applications and data can be replicated at different nodes so that users can be served by any of the nodes depending on the current load. This arrangement also provides high-availability since, in the event of a node failure, other nodes can still do the work. However, managing data replication in the ASP context is far more difficult than in Web search engines since applications can be update-intensive and both applications and databases must remain autonomous.

To obtain high-performance and high-availability, we replicate databases (and DBMSs) at several nodes, so they can be accessed in parallel through applications. Then the main problem is to assure the consistency of autonomous replicated databases. An obvious solution is synchronous replication which updates all replicas within the same (distributed) transaction. Synchronous replication enforces mutual consistency of replicas. However, it cannot scale up to large cluster configurations because it makes use of distributed transactions. A better solution that scales up is multi-master lazy replication. It provides high-availability and high-performance since replicas can be updated in parallel at different nodes. However, conflicting updates at different nodes can introduce replica divergence.

To address this problem, we have proposed a new preventive replication strategy that can avoid the occurrence of conflicts, by exploiting the cluster's high speed network, thus providing strong consistency, without the constraints of synchronous replication. We implemented a preventive replication manager, called RepDB*, on top of various DBMS (Oracle™, PostGresQL and BerkeleyDB). Our experimental results with an 8-node Linux cluster show that it scales-up and introduces a negligible loss of data freshness.

6.4.2. Distributed data processing

Participants: Esther Pacitti, Patrick Valduriez.

In a large-scale distributed system, at least two kinds of data access are difficult to support: update transactions which can be very frequent and deal with replicated data, and retrieval queries that deal with very large objects. The traditional distributed query processing strategies devise statically an optimal execution plan based on a cost model and statistics on the data, and execute it on selected nodes. This approach is not suited to our context where the node load can change rapidly and the cost of processing large objects is difficult to predict. Thus our approach is to devise new dynamic techniques.

In the context of a cluster system with update-intensive autonomous databases, we proposed a solution for a more dynamic load balancing. To increase parallel data processing, we use optimistic database replication with freshness control. First, we proposed algorithms to evaluate data freshness and compute the minimum set of refresh transactions needed to guarantee that a node is fresh enough with respect to a given query. Second, we proposed a solution to transaction routing that preserves database and application autonomy and a cost model to estimate replica freshness. We implemented our solution on a Linux cluster running Oracle™8i and performed extensive performance experiments using the TPC-C OLTP benchmark. Our results show that our solution outperforms existing solutions for typical transaction workloads. This work is done in cooperation with LIP6 in the context of the Leg@net RNTL project.

Regarding the problem of accessing large size objects, we proposed a novel approach, called Cherry Picking (CP), which is based on the modeling of data dependencies among expensive predicate input values as a k -partite graph. We showed how CP can be easily integrated into a cost-based query processor and demonstrated excellent performance results compared with static strategies. This work is done in cooperation with PUC-Rio, Brazil.

7. Contracts and Grants with Industry

7.1. Leg@net RNTL Project

The project involved Prologue Software, ASPLine and LIP6 in 2002-2003. The project leader for LIP6 is P. Valduriez (who started the project while he was a professor at Univ. Paris 6). The objective is to demonstrate the added value of the Application Service Provider model for managing autonomous applications and databases for pharmacy applications in France. To obtain a good performance/cost, the Leg@net project exploits cluster architectures, using data replication and load balancing, and Open Source Software, in particular Linux. The Atlas team has contributed to the data management system architecture and developed a preventive replication technique.

7.2. Microsoft Research Contract

The objective of the contract (2003-2006) is to contribute to the development of the ATL model management framework and foster the dissemination of our results as Open Source Software under a non restrictive license. In particular, we will port the ATL transformation components for metadata management, primarily developed for the Java platform, to the Microsoft .Net platform with use of COM components.

7.3. Domus Videum RNTL Project

The project (2002–2004) involves SFRS, INA, LaBRI, IRCCyN, IRISA (INRIA-Metiss/Vista/Temics) and IRIN (INRIA-Atlas) and is led by Thomson Multimedia. The main goal of this project is to provide services for a new generation of home multimedia platforms and more specifically, to develop an intelligent Personal Video Recorder (PVR). This kind of PVR is able to filter TV broadcasts depending on a user profile, and offer, through a user-friendly navigation interface, different points of view over programs (full version, summary, digest, best-of). The contribution of the Atlas team to the project is the task of learning the user profile based on metadata of TV programs, and the filtering process.

8. Other Grants and Activities

8.1. Regional Actions

The Atlas group participates in the COM project funded by the “Region des Pays de la Loire” (2000-2006). The objective of the COM project is to promote research in computer science in the region, in particular the creation of LINA (Laboratoire d’Informatique de Nantes Atlantique), a UMR between CNRS, University of Nantes and École des Mines de Nantes. J. Bézivin participates with the Obasco group in the “Club Objet de l’Ouest” which fosters cooperation in object technologies between public research laboratories (IRISA, University of Nantes, ENSTB, etc.), and industry (France Telecom, AQL, Softeam, etc.).

8.2. National Actions

We are involved in several CNRS specific actions:

- Databases and mobility (2002-2003) ;
- Personnalization of information (2003) ;

- Information retrieval: scaling up the size of corpus (2003);
- Model management (2003).

The Atlas group leads the MDP2P (Massive data management in peer-to-peer systems) project funded by the ACI “Masses of Data” of the French ministry of research. The MDP2P project is scheduled for 3 years (2003-2006) and involves three other INRIA groups: Paris and Texmex in Rennes, and Gemo in Orsay. The Atlas group is also leading the model engineering group at OFTA (Observatoire Français des Techniques Avancées)

8.3. International actions

We are involved in the following international actions :

- the Daad (Distributed computing with Autonomous Applications and Databases) project, partially funded by CNPQ and INRIA, with the universities PUC-Rio and UFRJ, Brazil, on distributed data management;
- the STIC multimedia network between France and Morocco, with University Mohammed V of Rabat, EMI, ENSIAS and University of Fès;
- the STIC Software Engineering project between France and Morocco with University Mohammed V of Rabat, EMI, ENSIAS and University of Fes;
- the OMG consortium, in which J. Bézivin participates to the MDA work.

Furthermore, we have regular scientific relationships with research laboratories in

- North America: Univ. of Waterloo (Tamer Özsu), University of California Berkeley (Michael Franklin), MIT (Stuart Madnick), New Jersey Institute of Technology (Vincent Oria), Wayne State University (Farshad Foutouhi and Wiliam Grosky), Kettering University (Peter Stanchev);
- Europe: CWI (Martin Kersten), University of Twente (Mehmet Aksit), University of Roskilde (Henrick Larsen);
- Others: University Federal of Rio de Janeiro (Marta Mattoso), Tokyo Metropolitan University (Hiroshi Ishikawa)

9. Dissemination

9.1. Animation of the scientific community

The members of the Atlas group have always been strongly involved in organizing the French database research community, in the context of the I3 GDR and of the conference Bases de Données Avancées (BDA). J. Bézivin is a member of the ALP GDR and is involved in the French object research community, through the two main conferences Objet'XX and OCM'XX.

P. Valduriez is a member of the scientific committee of the ACI GRID and a member of the ACM SIGMOD steering committee. In 2004, the ACM SIGMOD conference will be held in Paris (for the first time outside North America); P. Valduriez is general chair and J. Martinez is web site manager.

9.2. Editorial Program committees

Participation in the editorial board of scientific journals:

- Int. Journal on Intelligent and Cooperative Database Systems, World Scientific: P. Valduriez
- Distributed and Parallel Database Systems, Kluwer Academic Publishers: P. Valduriez

- Internet and Databases: Web Information Systems, Kluwer Academic Publishers: P. Valduriez
- Ingenierie des Systemes d'Information, Hermes : N. Mouaddib, P. Valduriez
- Journal of Object Technology: J. Bézivin
- SoSyM, Software and System Modeling, Springer Verlag: J. Bézivin

Participation in conference program committees :

- ACM SIGMOD Int. Conf. on the Management of Data (SIGMOD), 2004: E. Pacitti
- ACM Int. Symp. on Applied Computing: Multimedia and Visualisation Track (SAC), 2004: J. Martinez
- Int. Conf. on Very Large Databases (VLDB), 2003: P. Valduriez
- Int. Conf. on Web Information System Engineering (WISE), 2003: P. Valduriez
- Int. Data Engineering and Applications Symposium (IDEAS), 2003: E. Pacitti
- Brazilian Symposium on Databases (SBBD), 2003: E. Pacitti, P. Valduriez
- Int. Conf. on Enterprise Information Systems (ICEIS), 2003: J. Bézivin
- Int. Conf. on the Unified Modeling Language (UML), 2003: J. Bézivin
- Journées Recherche d'Information Assistée par Ordinateur (RIAO), 2004: J. Martinez
- Conference on Recherche d'Information et Applications (CORIA), 2004: J. Martinez
- Journées Bases de Données Avancées (BDA), 2003: J. Martinez, E. Pacitti, P. Valduriez
- Joint Workshop "Information retrieval: scaling up the size of corpus" with Congrès sur l'Informatique des Systèmes d'Information et de Décision (INFORSID), 2003: J. Martinez

9.3. Invited Talks

P. Valduriez gave an invited talk at the MIAGE days at IFSIC, Rennes, and at the RNTL days, Grenoble. He will be a keynote speaker at the Vecpar 2004 conference in Valence.

J. Bézivin gave a tutorial at the summer school EDF/CEA/INRIA. He will be a keynote speaker at the UML'2003 conference in San Francisco.

9.4. Teaching

All the members of the Atlas group teach database management, multimedia, and software engineering at the Bs, Ms and Ph.D. degree level at the University of Nantes.

The book *Principles of Distributed Database Systems*, co-authored with professor Tamer Özsu, U. Waterloo, published by Prentice Hall in 1991 et 1999 (2nd edition) has become the standard book for teaching distributed databases all over the world. Our Web site features course material, exercises, and direct communication with professors.

10. Bibliography

Major publications by the team in recent years

- [1] T. ÖZSU, P. VALDURIEZ. *Principles of Distributed Database Systems, 2nd edition*. Prentice Hall, 1999.
- [2] J. BÉZIVIN, N. PLOQUIN. *Tooling the MDA framework: a new software maintenance and evolution scheme proposal*. in « Journal of Object-Oriented Programming, JOOP », 2001.

- [3] L. BOUGANIM, D. FLORESCU, P. VALDURIEZ. *Load Balancing for Parallel Query Execution on NUMA Multiprocesseurs*. in « Distributed and Parallel Databases », number 1, volume 7, 1999.
- [4] M. GELGON, P. BOUTHEMY. *A region-level motion-based graph representation and labeling for tracking a spatial image partition*. in « Pattern Recognition », number 4, volume 33, 2000.
- [5] J. MARTINEZ. *Database Browsing*. in « Journal of Research and Practice in Information Technology (JRPIT), Australian Computer Society Inc. », number 1, volume 34, 2002.
- [6] J. MARTINEZ, N. MOUADDIB. *Multimedia and Databases: A Survey*. in « Networking and Information Systems Journal (NISJ) », number 1, volume 2, 1999.
- [7] E. PACITTI, P. MINET, E. SIMON. *Replica Consistency in Lazy Master Replicated Databases*. in « Distributed and Parallel Databases », number 3, volume 9, 2001.
- [8] E. PACITTI, E. SIMON. *Update Propagation Strategies to Improve Freshness in Lazy Master Replicated Databases*. in « The VLDB Journal », number 3-4, volume 8, 2000.
- [9] P. PUCHERAL, L. BOUGANIM, P. VALDURIEZ, C. BOBINEAU. *PicoDBMS : Scaling down database techniques for the Smartcard*. in « The VLDB Journal, special issue on Best Papers from VLDB2000 », number 2-3, volume 10, 2001.
- [10] G. RASCHIA, N. MOUADDIB. *A Fuzzy Set-based Approach to Database Summarization*. in « Int. Journal of Fuzzy Sets and Systems », number 2, volume 129, 2002.

Doctoral dissertations and “Habilitation” theses

- [11] R. LUTFI. *Indexation “intelligente” et recherche par le contenu de l’audio*. Ph. D. Thesis, University of Nantes, 2003.
- [12] M. PELTIER. *Techniques de transformation de modèles basées sur la méta-modélisation*. Ph. D. Thesis, University of Nantes, 2003.

Articles in referred journals and book chapters

- [13] T. ÖZSU, P. VALDURIEZ. A. TUCKER, editor, *Handbook of Computer Science and Engineering, 2nd Edition*. CRC Press, 2004, chapter Distributed and Parallel Database Systems, to appear.
- [14] G. BERNARD, J. BEN-OTHTMAN, L. BOUGANIM, G. CANALS, B. DEFUDE, J. FERRIÉ, S. GANARSKI, R. GUERRAOU, P. MOLLI, P. PUCHERAL, C. RONCANCIO, P. SERRANO-ALVARADO, P. VALDURIEZ. *Mobilité et bases de données : état de l’art et perspectives*. in « Techniques et Sciences de l’Informatique (TSI) », number 3-4, volume 22, 2003.
- [15] G. BERNARD, J. BEN-OTHTMAN, L. BOUGANIM, G. CANALS, B. DEFUDE, J. FERRIÉ, S. GANARSKI, R. GUERRAOU, P. MOLLI, P. PUCHERAL, C. RONCANCIO, P. SERRANO-ALVARADO, P. VALDURIEZ. *Mobile Databases: a report on open issues and research directions*. in « ACM SIGMOD Record », 2004, to appear.

- [16] H. NAACKE, F. D. NGOC, P. VALDURIEZ. *Routage de Transactions dans un Cluster de Bases de Données Autonomes*. in « Ingénierie des Systèmes d'Information (ISI) », 2004, to appear.
- [17] A. PIGEAU, M. GELGON. *Organizing a personal image collection with statistical model-based ICL clustering on spatio-temporal camera phone meta-data*. in « Journal of Visual Communication and Image Representation », 2004, to appear.
- [18] R. SAINT-PAUL, G. RASCHIA, M. MOUADDIB. *Résumé de bases de données : application au domaine bancaire*. in « Technique et Science Informatiques (TSI) », 2003.

Publications in Conferences and Workshops

- [19] J. BÉZIVIN, G. DUPÉ, F. JOUAULT, G. PITETTE, J. ROUGUI. *First Experiments with the ATL Transformation Language: transforming XSLT into Xquery*. in « OOPSLA Workshop », 2003.
- [20] J. BÉZIVIN, N. FARCET, J.-M. JÉZÉQUEL, B. LANGLOIS, D. POLLET. *Reflective Model Driven Engineering*. in « Int. Conf. on UML », 2003.
- [21] J. BÉZIVIN, S. GÉRARD, P.-A. MULLER. *MDA components: Challenges and Opportunities*. in « Int. Workshop on Metamodelling for MDA », 2003.
- [22] J. BÉZIVIN, P.-A. MULLER, P. STUDER. *Platform Independent Web application Modelling*. in « Int. Conf. on UML », 2003.
- [23] L. CARMINATI, J. BENOIS, M. GELGON. *Human detection and tracking for video surveillance applications in a low-density environment*. in « SPIE Visual Communications and Image Processing conference », 2003.
- [24] S. GANARSKI, C. LEPAPE, P. VALDURIEZ. *Relaxing Freshness for Improving Performance in a Cluster of Replicated Databases*. in « Int. Workshop on Distributed Data & Structures (WDAS) », 2003.
- [25] S. GANARSKI, C. LEPAPE, P. VALDURIEZ. *Trading Freshness for Performance in a Cluster of Replicated Databases*. in « Int. Conf. on Cooperative Information Systems (CoopIS) », 2003.
- [26] E. LOISANT, R. SAINT-PAUL, J. MARTINEZ, G. RASCHIA, N. MOUADDIB. *Browsing Clusters of Similar Images*. in « Actes des 19e Journées Bases de Données Avancées (BDA'2003) », 2003.
- [27] I. MBAYE, J. MARTINEZ, R. O. H. THAMI. *Un modèle d'indexation pour la vidéo*. in « Première Conférence Plénière STIC (CoPSTIC) », 2003.
- [28] I. MBAYE, J. MARTINEZ, R. O. H. THAMI. *Un modèle de découpage progressif et d'indexation de la vidéo par le contenu*. in « 2^e Journées NTIC », 2003.
- [29] H. NAACKE, F. D. NGOC, P. VALDURIEZ. *Transaction Routing with Freshness Control in a Cluster of Replicated Databases*. in « Actes des 19e Journées Bases de Données Avancées (BDA'2003) », 2003.

- [30] E. PACITTI, T. ÖZSU, C. COULON.. *Preventive Multi-Master Replication in a Cluster of Autonomous Databases*. in « Int. Conf. on Parallel and Distributed Computing (Euro-Par 2003) », 2003.
- [31] E. PACITTI, T. ÖZSU. *Multi-Master Replica Consistency in a Cluster of Peer-to-Peer Autonomous Databases*. in « Actes des 19e Journées Bases de Données Avancées (BDA'2003) », 2003.
- [32] A. PIGEAU, M. GELGON. *Spatio-temporal organization of one's personal image collection with model-based ICL-clustering*. in « 3rd Int. Workshop on Content-Based Multimedia Indexing (CBMI'2003) », 2003.
- [33] A. PIGEAU, M. GELGON. *Organisation statistique spatio-temporelle d'une collection d'images acquises d'un terminal mobile géolocalisé*. in « Congrès Reconnaissance des Formes et Intelligence Artificielle (RFIA'2004) », 2004, To appear.
- [34] A. PIGEAU, G. RASCHIA, M. GELGON, N. MOUADDIB, R. SAINT-PAUL. *A fuzzy linguistic summarization technique for TV recommender systems*. in « IEEE Int. Conf. of Fuzzy Systems », 2003.
- [35] F. PORTO, E. LABER, P. VALDURIEZ. *Cherry Picking: a semantic query processing strategy for the evaluation of expensive predicates*. in « Brazilian Symposium on Databases (SBBD) », 2003.

Internal Reports

- [36] BÉZIVIN, E. BRETON, P. VALDURIEZ, G. DUPÉ. *The ATL Transformation-Based Model Management Framework*. Technical report, IRIN, 2003.