



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team ATLAS

*Complex Data Management in Distributed
Systems*

Rennes

THEME SYM

Activity
R *eport*

2005

Table of contents

1. Team	1
2. Overall Objectives	1
2.1. Overall Objectives	1
3. Scientific Foundations	2
3.1. Scientific Foundations	2
3.1.1. Data Management	2
3.1.2. Multimedia Data Management	3
3.1.3. Fuzzy Logic	3
3.1.4. Model Engineering	4
3.1.5. Distributed Data Management	4
4. Application Domains	5
4.1. Application Domains	5
5. Software	6
5.1. ATL (Atlas Transformation Language)	6
5.2. AMW (Atlas Model Weaver)	6
5.3. $Find_{AGE}^{Im}$	7
5.4. RepDB*	7
5.5. SaintEtiQ	7
6. New Results	7
6.1. Database summaries	7
6.1.1. Optimization of database summarization	8
6.1.2. Querying the summaries: multidimensional indexing	9
6.1.3. On-Line Analytical Processing of summaries	9
6.1.4. Summaries over a P2P architecture	9
6.2. Model management	10
6.2.1. Model transformation	10
6.2.2. Model weaving applied to data mapping	11
6.2.3. Global model management	11
6.3. Multimedia data management	12
6.3.1. Personal image collection management from mobile devices	12
6.3.2. Decentralized, distributed learning of multimedia class models	12
6.3.3. Selecting a discriminant subset of texture features for DBMS indexing	13
6.4. Distributed data management	13
6.4.1. Optimizing preventive data replication in a database cluster	14
6.4.2. Freshness-aware transaction routing in a database cluster	14
6.4.3. OLAP query processing in a database cluster	15
6.4.4. Data management in APPA	15
7. Contracts and Grants with Industry	16
7.1. Microsoft Research (2003-2006)	16
7.2. IBM/Eclipse (2004-2005)	16
7.3. RNTL Modathèque (2004-2005)	16
7.4. Caroll Motor (2003-2006)	16
7.5. IP Modelware (2004-2006)	17
7.6. Ilog (2005)	17
8. Other Grants and Activities	17
8.1. Regional Actions	17
8.1.1. COM, Région Pays-de-la-Loire (2000-2006)	17

8.1.2.	GeoPict, Mégalis, Région Pays-de-la-Loire (2005-2006)	17
8.2.	National Actions	17
8.2.1.	ACI MDP2P (2003-2006)	17
8.2.2.	ACI SemWeb (2004-2007)	18
8.2.3.	ACI APMD (2004-2007)	18
8.3.	International actions	18
9.	Dissemination	19
9.1.	Animation of the scientific community	19
9.2.	Editorial Program committees	19
9.3.	Invited Talks	20
9.4.	Teaching	20
10.	Bibliography	20

1. Team

Head of Project Team

Patrick Valduriez [Director of Research, INRIA]

Vice-head of Project Team

Jean Bézivin [Professor, University of Nantes]

Faculty Members, University of Nantes

Marc Gelgon [Associate Professor]

José Martinez [Professor]

Noureddine Mouaddib [Professor]

Esther Pacitti [Associate Professor]

Guillaume Raschia [Associate Professor]

Technical Staff

Freddy Allilaire [Engineer, CDD Modelware]

David Touzet [Engineer, CDD Modelware]

Gaëtan Gaumer [Engineer, CDD MDP2P, until september]

Peter Rosenthal [Engineer, CDD Modelware, until march]

Administrative Assistant

Marie Pivaut [Secretary, CDD]

Ph.D. Students

Reza Akbarinia [Fellowship Iranian gov.]

Cédric Coulon [Fellowship INRIA until August, then ATER Univ. of Nantes]

Marcos Del Fabro Didonet [Fellowship Microsoft]

Rabab Hayek [Fellowship MENRT since September]

Frédéric Jouault [Fellowship INRIA]

Erwan Loisant [Fellowship Monbusho, Japan]

Jorge Manjarrez Sanchez [Fellowship Conacyt, Mexico]

Vidal Martins [Univ. PUCPR, Brazil]

Afshin Nikseresht [Fellowship Iranian gov.]

Antoine Pigeau [Fellowship MENRT]

Jorge Quiane Ruiz [Fellowship Conacyt, Mexico]

Régis Saint-Paul [Assistant prof. (ATER) until September]

Stéphane Tayeb [Fellowship MENRT]

Amenel Voglozin [Fellowship MENRT]

Post-doctoral fellow

Ivan Kurtev [Fellowship Modelware since June]

2. Overall Objectives

2.1. 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 (cluster systems, P2P 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 which we address in four themes. 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 personalised access to multimedia data. Finally, the theme “distributed data management” addresses the problems of data replication and distributed query processing with complex data.

These dimensions are not independent and we foster cross-fertilization between themes. Examples of recent inter-theme research activities are: multimedia database summaries, multimedia data management in cluster systems, database summaries in P2P systems, and model management applied to distributed data integration.

3. Scientific Foundations

3.1. Scientific Foundations

Keywords: *Data management, database, distributed database, distributed systems, fuzzy logic, model engineering, multimedia, summaries.*

3.1.1. Data Management

Data management is concerned with the storage, organisation, 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. The relational model, by resting on a strong theory (set theory and firsts-order logic) to provide data independence, has revolutionized database management. The major innovation of relational DBMS has been to allow data manipulation through queries expressed in a high-level (declarative) language such as SQL. Queries can then be automatically translated into optimized query plans that take advantage of underlying access methods and indices. Many other advanced capabilities have been made possible by data independence : data and metadata modelling, schema management, consistency through integrity triggers, transaction support, etc.

This data independence principle has also enabled DBMS to continuously integrate new advanced capabilities such as objet and XML support and to adapt to all kinds of hardware/software platforms from very small smart devices (PDA, smart card, etc.) to very large computers (multiprocessor, cluster, etc.) in distributed environments.

Following the invention of the relational model, research in data management continued with the elaboration of strong database theory (query languages, schema normalization, complexity of data management algorithms, transaction theory, etc.) and the design and implementation of DBMS. For a long time, the focus was on providing advanced database capabilities with good performance, for both transaction processing and decision support applications. And the main objective was to support all these capabilities within a single DBMS.

Today’s hard problems in data management go well beyond the traditional context of DBMS. These problems stem from the need to deal with data of all kinds, in particular, text and multimedia, in highly

distributed environments. Thus, we also capitalize on scientific foundations in multimedia data management, fuzzy logic, model engineering and distributed systems to address these problems.

3.1.2. Multimedia Data Management

Multimedia data such as image, audio or video is quite different from structured data and semi-structured (text) data in that it is media-specific (with specific operations) and described by metadata. Furthermore, useful representations of multimedia data, that are involved in storage and computation phases, are possibly voluminous and generally defined in high-dimensional spaces. Multimedia data management aims at providing high-level capabilities for organizing, searching and manipulating multimedia collections efficiently and accurately. To address this objective, we rely on the following research areas which we list in an order corresponding to the data flow: multimedia data analysis and pattern recognition, information retrieval and databases (mostly distributed). The overall architecture remains organised around the three fundamental parts of database design: modelling, querying and indexing. However, they have to be considerably adapted in order to manipulate multimedia data while maintaining the desired abstraction level.

With respect to modelling, multimedia data analysis performs automatic translation of raw multimedia data into sets of discriminant, concise descriptions that are used for indexing and searching. These descriptions range from low-level transforms on the original data (e.g. image texture features), that translate into feature vectors, to more abstract representations (e.g. parametric models), that often attempt to capture a class rather than an instance of multimedia elements. Furthermore, media content creators may add metadata information that conveys more semantics. Briefly stated, multimedia data analysis deals with the design of suitable observations from multimedia and pattern recognition techniques. Its interdependence with information retrieval and databases has encouraged the development of dedicated research branches, since many interesting applications consider multimedia information retrieval on voluminous data. Our work follows this direction.

Querying has been concerned with the conceptual access to data by the user with a high-level (SQL-like) query language on user-defined schemas. In contrast, techniques for querying multimedia data come from the information retrieval community. Although extensible, each content-based multimedia system relies on a single, well-defined schema (similar to the document-term matrix from textual documents). Similarly, the common query in multimedia is a similarity search where the objects retrieved are ordered according to some scores based on a distance function defined on a feature vector, rather than a boolean expression. Similarly, relevance feedback has been introduced early in content-based systems since it is impossible to provide a concise description of a user's needs. In this respect, multimedia querying becomes mainly an interactive activity. Finally, it appears that several difficulties can be overcome by clustering multimedia data, something which is not new in databases, e.g., datawarehouses, but has to be done in a totally different way.

These important differences lead to reconsidering indexing too. Indexing is concerned with the physical access to multimedia data. The aim of indices is to rapidly access the data requested by the query. Efficient multimedia descriptors often span high dimensional spaces (say, 10 to 1,000 dimensions) since, to some extent, more features means more discriminant. Application of classical indexing structures (tree-based and hashing-based) supplied by database research is not effective, at least not in the straightforward manner, because these structures suffer from the "dimensionality curse problem", which states that the performance of indexing (and thus querying) degrades *severely* as the data dimensionality increases, in particular in the abovementioned dimension range. This particular issue is currently attracting much interest. The general problem is to achieve both high *effectiveness*, i.e., retrieving multimedia data that correspond to the user's needs and *efficiency* in order to scale up to large multimedia databases.

3.1.3. Fuzzy Logic

The ever growing size of databases makes data summarization needed in order to present the user a concise, yet complete view of the data. Our proposed summarization process [58] can roughly be described as a two step process. The first step is to rewrite the original database data into an homogeneous user-oriented vocabulary. The second step is then to use a concept formation algorithm against the rewritten data. The fuzzy-set theory provides a mathematical foundation to handle these two steps in a more efficient and robust

way than can be achieved with first order logic. Fuzzy sets theory was introduced by L.A. Zadeh in 1965 in order to model sets whose boundaries are not sharp. A fuzzy (sub)set F of an universe Ω is defined thanks to a membership function denoted by μ_F which maps every element x of Ω into a degree $\mu_F(x)$ in the unit interval $[0, 1]$. Thus, a fuzzy set is a generalization of regular set (whose membership function is defined on the pair $(0,1)$).

In the first step, database tuples are rewritten using a user defined vocabulary. This vocabulary is intended to match as well as possible the natural language in which users expresses their knowledge. A database user usually refers to his or her data using a vocabulary appropriate for his field of expertise and understood by his or her fellows. For example, a salary will be said to be high, reasonable or average. This description in fact is an implicit categorization and there is no crisp border line between an average and a high salary. Fuzzy logic offers the mathematical ground to define such a vocabulary in term of linguistic variables where each data is more or less satisfactorily described by the concept.

In a concept formation algorithm, new data are incorporated into a concept hierarchy using a local optimization criteria to decide how should the hierarchy be modified. A quality measure is evaluated to compare the effect of operators that modify the hierarchy topology namely, creating a new node, creating a new level, merging two nodes, or splitting one. Using fuzzy logic in the evaluation of this measure, our concept formation algorithm is less prone to suffer the well known threshold effect of similar incremental algorithm.

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 handle this operation. Using the database summary, queries with too few results can be relaxed to retrieve partially satisfactory subsets of the database. The fuzzy matching mechanism also permit to handle user queries expressed in vague or imprecise terms.

3.1.4. Model Engineering

A model is a formal description of a design artifact such as a relational schema, an XML schema, a UML model or an ontology. Data and meta-data modelling have been studied by the database community for a long time. We also witness the impact of similar principles in software engineering. Metamodels are used today to define domain specific languages that may help capturing the various aspects of complex 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 and, as such, is leverage for synchronization. Modelling naturally comes along with correspondences and constraints between models, i.e. the representation of a system by a model, the conformance of a model to a metamodel and the relation of one metamodel with another expressed by a transformation. In this area, research focuses on constraint languages and the traceability of transformations.

Considering models, meta-models, and model transformations as first class elements yields much genericity and flexibility to build complex data-intensive systems. A central problem of these systems is data mapping, i.e. mapping heterogeneous data from one representation to another. Examples can be found in different contexts such as schema integration in distributed databases, data transformation for data warehousing, data integration in mediator systems, data migration from legacy systems, ontology merging, schema mapping in P2P systems, etc. A data mapping typically specifies how data from one source representation (e.g. a relational schema) can be translated to a target representation (e.g. another, different relational schema or an XML schema). Generic model management has recently gained much interest to support arbitrary mappings between different representation languages.

3.1.5. Distributed Data Management

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 capitalise 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.

Data management in distributed systems has been traditionally achieved by distributed database systems which enable users to transparently access and update several databases in a network using a high-level query language (e.g. SQL) [2]. Transparency is achieved through a global schema which hides the local databases' heterogeneity. In its simplest form, a distributed database system is a centralized server that supports a global schema and implements distributed database techniques (query processing, transaction management, consistency management, etc.). This approach has proved effective for applications that can benefit from centralized control and full-fledge database capabilities, e.g. information systems. However, it cannot scale up to more than tens of databases. Data integration systems extend the distributed database approach to access data sources on the Internet with a simpler query language in read-only mode.

Parallel database systems also extend the distributed database approach to improve performance (transaction throughput or query response time) by exploiting database partitioning using a multiprocessor or cluster system. Although data integration systems and parallel database systems can scale up to hundreds of data sources or database partitions, they still rely on a centralized global schema and strong assumptions about the network.

In contrast, peer-to-peer (P2P) systems adopt a completely decentralized approach to data sharing. By distributing data storage and processing across autonomous peers in the network, they can scale without the need for powerful servers. Popular examples of P2P systems such as Gnutella and Kaaza have millions of users sharing petabytes of data over the Internet. Although very useful, these systems are quite simple (e.g. file sharing), support limited functions (e.g. keyword search) and use simple techniques (e.g. resource location by flooding) which have performance problems. To deal with the dynamic behavior of peers that can join and leave the system at any time, they rely on the fact that popular data get massively duplicated.

Initial research on P2P systems has focused on improving the performance of query routing in the unstructured systems which rely on flooding. This work led to structured solutions based on distributed hash tables (DHT), e.g. CAN and CHORD, or hybrid solutions with super-peers that index subsets of peers. Although these designs can give better performance guarantees, more research is needed to understand their trade-offs between fault-tolerance, scalability, self-organization, etc.

Recently, other work has concentrated on supporting advanced applications which must deal with semantically rich data (e.g., XML documents, relational tables, etc.) using a high-level SQL-like query language. Such data management in P2P systems is quite challenging because of the scale of the network and the autonomy and unreliable nature of peers. Most techniques designed for distributed database systems which statically exploit schema and network information no longer apply. New techniques are needed which should be decentralized, dynamic and self-adaptive.

4. Application Domains

4.1. Application Domains

Keywords: *Application Service Provider (ASP), distributed collaborative application, 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-scale distributed collaborative applications, 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 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 by exploiting data replication and load balancing techniques.

Large scale distributed collaborative applications are getting common as a result of the progress of distributed technologies (GRID, P2P, and mobile computing). Consider a professional community whose members wish to elaborate, improve and maintain an on-line virtual document, e.g. reading or writing notes on classical literature, or common bibliography, supported by a P2P system. They should be able to read/write on the application data. An important aspect of large scale distributed collaborative applications is that user nodes may join and leave the network whenever they wish, thus hurting data availability. In Atlas, we address the issues of replication, query processing and load balancing for such applications assuming a P2P architecture (APPA) that is fully decentralized.

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 audio-visual material stored either locally on a PC/Settop-box, or a mobile handset. Such domestic applications, extended to the video medium mainly, appear as a natural perspective for future 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.

URL: <http://www.eclipse.org/gmt/>

ATL is a transformation-based model management framework, with metadata management and data mappings as the main applications. The ATL language is designed to be general and abstract. We 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 implemented in Java, and we are porting major transformation components to the .Net platform. ATL has been registered in 2004 (together with TNI-Software and the University of Nantes) to the APP (Agence pour la Protection des Programmes) and is released as Open Source Software under the Eclipse Public Licence.

5.2. AMW (Atlas Model Weaver)

Participants: Jean Bézivin, Marcos Didonet Del Fabro, Patrick Valduriez.

URL: <http://www.eclipse.org/gmt/>

AMW is a component-based platform for model weaving, i.e. establishing and managing correspondences between models. The platform is based on the Eclipse contribution mechanism: components are defined in separated plugins. The plugins are further interconnected to create the model weaver workbench. Components for user interface, matching algorithms and serialization of models may be plugged as necessary. We extended

the Eclipse EMF architecture for model manipulation to coordinate the weaving actions. We use the EMF reflective API to obtain a standard weaving editor which adapts its interface according to metamodels modifications. The ATL transformation engine is plugged as the standard transformation platform. ATL has been registered in 2005 (together with the University of Nantes) to the APP (Agence pour la Protection des Programmes) and is released as Open Source Software under the Eclipse Public Licence.

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

Participants: José Martinez, Erwan Loisant.

$Find_{AGE}^{Im}$ 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.

5.4. RepDB*

Participants: Cédric Coulon, Gaëtan Gaumer, Esther Pacitti, Patrick Valduriez.

URL: <http://www.sciences.univ-nantes.fr/lina/ATLAS/RepDB/>

RepDB* is a data management component for replicating autonomous databases or data sources in a cluster system. It has been initially designed in the context of the Leg@net RNTL project and further developed in the context of the ACI MDP2P project. RepDB* supports preventive data replication capabilities (multi-master modes, partial replication, strong consistency) 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 64-node cluster at INRIA-Rennes. In 2004, we registered RepDB* (together with the University of Nantes) to the APP (Agence pour la Protection des Programmes) and released it as Open Source Software under the GPL licence.

5.5. SaintEtiQ

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

URL: <http://www.simulation.fr/seq>

SAINTETIQ is a data summarisation system which provides 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 summarisation process. Finally, a complete graphical user interface has been developed to support the user manipulating and browsing data, background knowledges and summaries. SAINTETIQ is now available as a Web service.

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. Our recent work has focused on optimisation of the overall architecture for database summarization. We also dealt with decision support system based on summaries rather than datacubes. Finally, we carried further our work on summary querying techniques. These directions provide different ways of handling summaries coupled with a DBMS.

6.1.1. Optimization of database summarization

Participants: Nouredine Mouaddib, Guillaume Raschia, Régis Saint-Paul.

Our summarisation process is intended to seamlessly become part of existing DBMS, and then, requires low memory consumption, a reliable serialization system and effective linear time complexity for atomic operations. To this end, we proposed a service-oriented architecture for SAINTETIQ with the following main features [58], [59], [57]:

- models provided as XML Schemas for each input/output document (raw data, cooked data, summaries);
- stored procedures and triggers for Oracle and MS-SQL Server coupling;
- cache manager designed to maintain frequent summaries in memory and to serialize as binary streams the other ones;
- rewriting and summarization tasks are independently available as web services
- distributed computing facilities based on the implementation of each summary as an autonomous agent;

Cooked data are documents that contain a rewritten form of tuples and are the input of the summarization service. They provide a common representation of the data in the form of a set of descriptors for each attribute. Each descriptor is associated to a satisfaction degree. This makes the summarisation service independent of the underlying data type.

The summarisation service performs the learning task. It takes cooked data as input, and outputs a collection of summaries hierarchically arranged according to their precision. This hierarchical structure can later be used to produce a reduced version of any given size of the original relation R . New data are first incorporated in the root node of the hierarchy one at a time. Then, in a top-down clustering approach, data are processed from the root to the leaves. At each node, a measure evaluates the quality of hypothetical summary arrangements as resulting from applying a set of learning operators which locally modify the underlying partition, i.e. the set of children nodes.

The swap system supported by cache management drastically reduces memory consumption to a user or system-defined threshold. The SAINTETIQ cache policy disposes instances according to their use frequency. With this method, complete branches of the hierarchy are resident in memory while barely accessed summaries remain on the disk.

Our implementation is based on the Message-Oriented Programming paradigm. Each sub-system is autonomous and collaborates with the others through disconnected asynchronous method invocations. It is among the least demanding approaches in terms of availability and centralization. The autonomy of summary components allows for distributed computing. Once a component completes the treatment and evaluates the best operator for the hierarchy modification, if needed, a similar method is successively called on children nodes. The cache manager is able to handle several lists of summaries residing on different computers. The manager is also responsible for load balancing of the newly created process. Load balancing is achieved by analysing the number of processed tuples for each node: when data are treated in no particular order, the relative content size of a summary is a direct function of its frequency of use.

6.1.2. Querying the summaries: multidimensional indexing

Participants: Noureddine Mouaddib, Guillaume Raschia, Amenel Voglozin.

We investigated the area of multidimensional indexing from the point of view of space-partitioning. Through its architectural aspects, a summary hierarchy shares many features of the multidimensional indexes (R-Tree, UB-Tree, X-Tree, ...). Current work on flexible querying uses the hierarchy as an index to quickly select the appropriate database records, since in multidimensional indexing, each selection criterion reduces the search space for the other criteria [62], [20].

We proposed a querying mechanism for users to efficiently exploit the hierarchical summaries produced by SAINTÉTIQ. The first idea is to query the summaries using the vocabulary of the summarization process and taking advantage of the hierarchical structure of the summaries. The querying process answers queries in which the criteria specify labels from the vocabulary. The algorithms perform boolean set comparisons and use the tree structure to cut branches and quickly reduce the search space. This leads to important gains in response time, especially in the case of null answers (i.e., of an empty result set), as only a small part of the summaries hierarchy has to be explored, instead of the entire database.

Querying the summaries is meaningful to rapidly get a rough idea of the properties of tuples in a relation. However, queries may have a null answer. Apart from the cases where null answers are unacceptable, the user has to think about another query, which might fail as well, and so on. Thus, the second idea is to find ways to provide an approximate answer when the user's query produces an empty result. The intention, in an attempt to repair queries, is to offer an answer even when no summary match the query. We then proposed to modify the original query. This modification is performed from the optimistic idea that there exist results semantically close to those targeted by the user. In order to select such approximate-result summaries, the query is modified using the parsed SAINTÉTIQ summaries or pre-established information.

The next major step is the integration of this multidimensional index into a DBMS to assess the feasibility and performance issues.

6.1.3. On-Line Analytical Processing of summaries

Participants: Noureddine Mouaddib, Guillaume Raschia.

We proposed a general framework to explore and analyze database summaries built from massive data sets [51]. Summaries are self-descriptive and higher-level views of groups of raw data. We defined a logical data model called *summary partition*, as datacubes do in OLAP systems, in order to provide the end-user with a reduced and well-suited presentation of the data. Pre-built and ordered partitions are considered on the basis of a process dedicated to the generation of summaries at different levels of granularity. The overall on-line summarization processing is then intended to support a new approach to OLAP. To achieve this, we introduced a collection of algebraic operators over the space of summary partitions: relational, granularity and structuring operators are designed for on-line analytical processing of summarized versions of the data.

This work, yet preliminary, is intended to define the core algebra of an effective and rich tool for visualizing, querying and accessing summaries considered as compressed semantic views of raw data.

6.1.4. Summaries over a P2P architecture

Participants: Rabab Hayek, Noureddine Mouaddib, Guillaume Raschia, Patrick Valduriez.

We started to study the integration of a summary service into a P2P architecture in two directions [65]. The first direction is querying the summaries. Assume that we succeed in providing a global summary of all the data sources available on the P2P network. A peer submitting a query may get approximative information about data of other peers without having to visit them. Thus, the summary can serve as a global materialized view over the shared data. The second direction is querying data sources through summaries. In order to answer a query, a given peer must compute the set of relevant answering peers using available summaries. Hence, content-based peer targeting is more accurate than usual techniques based on structural information. It allows to reduce network traffic and to collect "meaningful answers".

As an application domain, we started to investigate one of the critical issues in web-based e-commerce. The problem is how to efficiently and effectively integrate and query heterogeneous e-catalogs. We proposed an integration framework for building and querying catalogs [11]. Our approach is based on a hybrid of peer-to-peer data sharing paradigm and webservices architecture. Peers in our system serve as domain-specific data integration mediators. Links between peers are established based on the similarity of the domain they represent. The relationships are used for routing queries among peers. As the number of catalogs involved increases, the need for filtering irrelevant data sources also increases. We applied a summarization technique to summarise the content of catalogs. The summaries are used to preselecting data sources that are relevant to a user query.

6.2. Model management

A model is a structure that represents a design artefact 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, Perl, 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 SQL to XQuery. There is much to gain if we could handle these various transformations in a more generic way with a coordinated family of languages. To contribute to this evolution, we have continued to refine ATL (Atlas Transformation Language) which is now the basis for a more general architecture named AMMA (ATLAS Model Management Architecture). Our research activities in AMMA concern model transformation and weaving, global management of related resources (mainly models, metamodels and transformations), and the integration of these functionalities into an open model management platform.

6.2.1. Model transformation

Participants: Freddy Allilaire, Jean Bézivin, Frédéric Jouault, Ivan Kurtev, David Touzet, Patrick Valduriez.

Model transformation, e.g. mapping a relational database schema into an XML schema, is a very useful and important operation in model management. We have proposed ATL, a combined declarative/imperative language that allows to transform source models into a target models. Like the source and target models, the transformation program is itself a model and thus conforms to a given metamodel. This corresponds to the general conceptual unifying scheme described in [12]. In addition to the language definition, a proof of concept has been proposed with first implementations on the Sun MDR/NetBeans environment and on Eclipse as part of the GMT open source project (eclipse.org/gmt). In addition to the ATL engine, a complete integrated development environment (IDE) has been built and also released as GMT open source. The IDE for ATL allows transformation editing and debugging (syntax coloring, step by step execution, breakpoints, environment observation, etc.). We developed several examples of model transformations as part of a basic library.

In the context of the ModelWare european integrated project, we collaborate with SINTEF, Norway, in applying ATL to several case studies. ATL is also being used by a strong international research community of more than 80 users and several companies including Airbus within the french project TOPCASED. It is aligned on the recent QVT normative recommendation [43]. However we have also shown how ALT allows to bridge the OMG and various other environments like GME or Microsoft DSL Tools [21], [24].

An original aspect of the ATL implementation is that it is based on the public definition of a portable transformation virtual machine. The specification of this virtual machine has been released on Eclipse. We

also proposed KM3 (Kernel MetaMetaModel), a domain specific language for specifying metamodels [23], [23] for example those describing tools' internal data formats (MS Excel, MS Project, MatLab, Bugzilla, Mantis, etc.).

Taking stock on the ATL implementation framework, we obtained several original results. We have shown the possibility to express a model verification by a pure transformation [28]. Not only the verification criteria may be expressed by a separate model, but also the diagnostic result may be expressed as a model conforming to a variable metamodel. We have proposed a proof of concept based on our ATL implementation. We are currently extending it to measure models by using a similar model-based organization. We have experimentally proven that a model transformation may produce an additional traceability model that may be used to relate the target to the source model after the transformation [42].

6.2.2. Model weaving applied to data mapping

Participants: Marcos Didonet Del Fabro, Jean Bézivin, Frédéric Jouault, Patrick Valduriez.

Mapping between heterogeneous data is a central problem in many data-intensive applications. A typical data mapping specifies how data from one source representation (e.g. a relational schema) can be translated to a target representation (e.g. an XML schema). Although data mappings have been studied independently in different contexts, there are two main issues involved. The first one is to discover the correspondences between data elements that are semantically related in the source and target representations. This is called schema matching in schema integration and many techniques have been proposed to (partially) automate this task. After the correspondences have been established, the second issue is to produce operational mappings that can be executed to perform the translation. Operational mappings are typically declarative, e.g. view definitions or SQL-like queries. However, using one mapping language causes serious limitations and makes mapping management difficult.

In [36], [35] we have proposed a solution based on model weaving which can better control the trade-off between genericity, expressiveness, and efficiency of mappings. In other words, our objective is to support generic data mapping (as in other model management systems but with a different approach) while exploiting specific mapping languages and engines, such as XQuery, SQL or ATL. Our solution considers mappings as models and exploits specific mapping engines. We defined model weaving as a generic way to establish element correspondences. Weaving models may then be used by a model transformation language to translate source model(s) into target model(s). We validated our approach using the ATLAS Model Weaver (AMW) prototype on several application scenarios. Our experiments have shown that many different proposals may be unified by our model-based approach. Coupling a weaving facility (like AMW) with a transformation facility (such as ATL) gave us good efficiency and flexibility. We have illustrated the joint use of AMW and ATL (i.e. generating executable transformations from correspondences) by several practical projects [48], [37], [21].

We have also shown in [22] that the model weaving operation may help solving important problems in software engineering as well. For instance, the so called "Y"-shaped software development cycle to build a software system can be solved by a merge from a business model and a platform model.

6.2.3. Global model management

Participants: Freddy Allilaire, Jean Bézivin, Frédéric Jouault, David Touzet, Patrick Valduriez.

Within a model management environment, the main elements produced or consumed are models, metamodels, correspondences or transformations [22], [30]. However, in order to allow for the manipulation of other resources such as XML documents, database tables or flat files, collections of generic importers and exporters are needed. Special attention should be given to the global management of all these resources. These models are explicitly typed by their corresponding metamodels allowing to define the signature of each tool. A platform has the precise and updated knowledge about all the connected tools, each one being characterized by its signature [10].

In our approach, all the information about the components known to a given platform is stored in a specific model named "megamodel". A megamodel is a kind of model registry that stores reference and metadata information on all accessible resources, including relations between these resources. It allows us to build

a minimal and highly extensible infrastructure. In particular, this allows easy extension of a local platform towards a distributed platform such as a P2P system without significant modification of tool interoperability mechanisms. Furthermore the approach fits well within the general conceptual scheme developed for model management. Experimental validation is being done through the AM3 (ATLAS MegaModel Manager) tool to record and control the global relations between model components [29].

6.3. Multimedia data management

The ability to store multimedia information in digital form has spurred both the 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 personalisation 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 have been working on the following forms of multimedia data : image, audio, video, and geo-temporal metadata possibly attached to a document.

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 both industrial (e.g. Microsoft Memex, project) and academic communities (e.g. ACM Multimedia Carpe workshops, ACM Sigmod 2005 keynote address). In particular, mobile devices such as camera-equipped phones are an interesting case for content creation and retrieval. In this context, we have proposed in 2004 an unsupervised technique for organizing, in an incremental fashion, an image collection, based on time and geo-location metadata. The objective is to recover the natural spatial, temporal or spatio-temporal structure present in such a data set.

This year, we have extended this work by adding the ability to build and maintain hierarchical temporal and geographical structures [53], [8]. The overall approach is founded on a hierarchy of mixture models. The scheme that provides incrementality of the hierarchy consists in introducing new data top down in the hierarchy and reconsidering (re-estimating) only sub-sections of the tree that appear to have changed structure. We conducted a validation on several home-made data sets, but the technique remains to be tested on sets that are both larger and more realistic.

6.3.2. Decentralized, distributed learning of multimedia class models

Participants: Marc Gelgon, Afshin Nikseresht.

A major need of multimedia indexing and retrieval is related to characterisation of classes (defining observations, capturing variability...), as addressed by other Inria project-teams for various medias, e.g. Lear, Vista, Imedia, Texmex and Metiss. Our work takes the following viewpoint. Large amounts of labelled and unlabelled multimedia data (in the sense of supervised learning) is distributed on the Internet in a large-scale, decentralized manner. On the other side, multimedia pattern recognition is generally a computation and data-intensive task. The perspective of exploiting collectively this data in a decentralized, distributed manner is made more realistic, for instance by current research efforts (e.g. Lear and Vista) towards the ability to learn class models despite clutter, thereby easing database collection. Positive learning examples for detecting faces are provided by images containing faces in a complex background.

As a first step, we have tackled the problem of estimating collectively the parameters and complexity of a Gaussian mixture for modelling class-conditional probability density. We selected Gaussian mixture models because of their versatility and ubiquitous use in the literature for capturing audio, video and image-based class models. We consider a set of nodes and a certain amount of labeled training data on each node. Models can be estimated locally, by a classical technique. So far, our focus has not been on locating suitable nodes in a network, but on aggregating local models to improve class characterization. The main feature of our approach

is that, in order to keep network and computation load moderate, participant nodes only transmit mixture model parameters rather than multimedia feature vectors, since aggregation only requires model parameters. Besides, an estimation of the appropriate number of components is carried out regularly, to enable scaling up. The proposed scheme aggregates local mixture model estimates by determining a suitable combination of their components. In fact, this combination is obtained by optimizing a modified Kullback divergence between the aggregated model and the concatenation of local models, through an iterative scheme. We conducted a preliminary validation of the scheme in the example of speaker recognition, spreading models parameters by gossip [53], [54]. The paper [53] obtained an IBM/Microsoft student grant award.

6.3.3. *Selecting a discriminant subset of texture features for DBMS indexing*

Participants: Ervan Loisant, José Martinez.

We have been arguing that an image DBMS should preferably use a standard DBMS technology and demonstrated in previous experimental work that maintaining efficiency in multimedia query processing is very difficult. Among the requirements, we pointed out that (i) reducing the initial size of metadata is extremely important (it impacts the space requirements, hence the time complexity), (ii) materialising some redundant information helps to improve queries (especially when the derived information is obtained through a complex formulation, which the database optimiser cannot recognise), (iii) parallelism is unavoidable *in fine*, as an additional improvement when all the other techniques have been used, (iv) distributing images on the nodes of a network, based on their visual properties has to be considered [66], (v) we have to take advantage of multi-dimensional indexing techniques, even limited ones, whenever possible.

Here, we addressed mainly points (i) and subsequently (v) for adding texture descriptors to existing colour-based ones [40], [41], [39]. The use of texture proved its effectiveness and usefulness in pattern recognition and computer vision. Unfortunately, texture descriptions are extremely varied. Firstly, we selected grey-level co-occurrence matrices (COM), a statistical approach that gives consistent and generally good enough results. Then, the problem was to extract a minimal subset of COM descriptors among the few tenth that have been proposed in the literature: contrast, directionality, roughness, variance et co-variance, etc. The reasons are the following. First, we have to evaluate the performance of texture features extracted from COM for retrieving similar textures, i.e., how *discriminative* they are. Second, due to hard limitations of multi-dimensional indexing (still harder for standard DBMSs), the feature vector's size has to be kept *minimal*, in the order of only four up to, say, twelve for state-of-the-art high-dimensional indexing techniques. Third, they have to be as *uncorrelated* as possible in order to avoid biases when evaluating queries or building classifications. Fourth, we prefer descriptors that could be in some way related to human perception since we would like to translate them into linguistic variables for a subsequent stage of our work, namely using our summarisation tool.

The selection process was based on *a priori* selection, followed by an experiment conducted on well-known test collections of textured images in order to extract the descriptors that fit our needs. This selection work has led to the following set of texture descriptors: *energy*, *entropy*, *variance*, *correlation* and *direction*.

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 high data autonomy, heterogeneity, and dynamicity.

We are interested in database clusters and peer-to-peer (P2P) systems which are good examples of large-scale distributed systems of high practical interest. However, to yield general results, we strive to develop common algorithmic solutions with the right level of abstraction from the context. In 2005, we continued our work on data management in database clusters and P2P systems. We pursued three research actions on optimizing preventive data replication, freshness-aware transaction routing and OLAP query processing in database clusters. We also continued the design of Atlas Peer-to-Peer Architecture (APPA) with new techniques for distributed reconciliation of replicated data and top-k query processing.

6.4.1. *Optimizing preventive data replication in a database cluster*

Participants: Cédric Coulon, Esther Pacitti, Patrick Valduriez.

To obtain high-performance and high-availability in a database cluster, we replicate databases (and DBMSs) at several nodes, so they can be accessed in parallel through applications. Then one main problem is to assure the consistency of autonomous replicated databases.

Preventive replication is an asynchronous solution that enforces strong consistency. Instead of using atomic broadcast, as in synchronous group-based replication, preventive replication uses First-In First-Out (FIFO) reliable multicast which is a weaker constraint. It works as follows. Each incoming transaction is submitted, via a load balancer, to the best node of the cluster. Each transaction T is associated with a chronological timestamp value C , and is multicast to all other nodes where there is a replica. At each node, a delay time d is introduced before starting the execution of T . This delay corresponds to the upper bound of the time needed to multicast a message. When the delay expires, all transactions that may have committed before C are guaranteed to be received and executed before T , following the timestamp chronological order (i.e. total order). Hence, this approach prevents conflicts and enforces consistency.

However, our original proposal has two main limitations. First, it assumes that databases are fully replicated across all cluster nodes and thus propagates each transaction to each cluster node. This makes it unsuitable for supporting large databases and heavy workloads on large cluster configurations. Second, it has performance limitations since transactions are performed one after the other, and must endure waiting delays before starting. Thus, refreshment is a potential bottleneck, in particular, in the case of bursty workloads where the arrival rates of transactions are high at times.

In [33], we introduced support for partial replication, where databases are partially replicated at different nodes. Unlike full replication, partial replication can increase access locality and reduce the number of messages for propagating updates to replicas. In [31], we proposed several optimizations to eliminate delay times and allow for the concurrent execution of the transactions. In [19], we describe our complete solution with significant extensions regarding replication configurations, concurrency management, proofs of algorithms' correctness and performance evaluation. This work is done in cooperation with Tamer Özsu, University of Waterloo.

We implemented our algorithms in our RepDB* prototype running on top of the PostgreSQL Open Source DBMS and performed extensive experiments over the 64-node Linux cluster of the Paris project-team at IRISA. Our experimental results [60] using the TPC-C Benchmark have shown that they yield excellent scale-up and speed up.

6.4.2. *Freshness-aware transaction routing in a database cluster*

Participants: Esther Pacitti, Patrick Valduriez.

In a database cluster with replication, maintaining replicas mutually consistent may hurt performance. However, there are important cases where consistency can be relaxed. For instance, read-only queries do not always require reading perfectly consistent data and may tolerate inconsistencies. Thus, an interesting solution is to trade consistency for performance based on users' requirements. In most approaches (including ours), consistency reduces to freshness: update transactions are globally serialised over the different cluster nodes, so that whenever a query is sent to a given node, it reads a consistent state of the database. Global consistency can be achieved using either a preventive approach that avoids conflicts such as [19], or an optimistic approach with conflict detection and reconciliation. However, the consistent state may not be the latest one, since update transactions may be running at other nodes. Then the data freshness of a node reflects the divergence between its actual state and the state it would have if all transactions had already been applied to it.

In [16], we proposed a freshness model which allows users to specify freshness requirements for their queries. This model allows defining conflict classes between queries and transactions. In [13], using this freshness model, we described the design and implementation of the Leganet system which performs freshness-aware transaction routing in a database cluster. We use multi-master replication and relaxed replica freshness

to increase load balancing. The Leganet system preserves database and application autonomy using non intrusive techniques that work independently of any DBMS. The main contribution is a transaction router which takes into account freshness requirements of queries at the relation level to improve load balancing. This router uses a cost function that takes into account not only the cluster load in terms of concurrently executing transactions and queries, but also the estimated time to refresh replicas to the level required by incoming queries. Using the Leganet prototype implemented at LIP6 on an 11-node cluster running Oracle8i and using emulation up to 128 nodes, our validation based on the TPC-C benchmark has demonstrated the performance benefits of our approach.

In most approaches to load balancing including Leganet, refreshment is tightly-coupled with other functions such as scheduling and routing. This makes it difficult to analyze the impact of the refresh strategy itself. Many refresh strategies have been proposed in the context of distributed databases, data warehouse and database clusters. In [46], we proposed a model which allows describing and analyzing existing refresh strategies, independent of load balancing issues. We described an experimental validation based on a workload generator, to test some typical strategies against different workloads. The results show that the choice of the best strategy depends not only on the workload itself, but also on the conflict rate between transactions and queries and on the level of freshness required by queries. Although there is no strategy that is best in all cases, we found that one strategy, As Soon As Underloaded, is usually very good and could be used as default strategy. This work is done in cooperation with LIP6.

6.4.3. OLAP query processing in a database cluster

Participants: Esther Pacitti, Patrick Valduriez.

OLAP applications require high-performance database support. They typically access large data sets using heavy-weight read-intensive queries. A simple, yet efficient, solution to OLAP query processing in database clusters is virtual partitioning which yields intra-query parallelism. The main requirement for virtual partitioning to work is that each node has access to the entire database. This can be achieved either in a shared-disk architecture or in a shared-nothing architecture with full replication. Each query submitted to the database cluster is rewritten in a number of queries, one for each node, by adding range predicates that correspond to different virtual partitions of a relation.

The main advantages of virtual partitioning are flexibility for node allocation during query processing tasks and high cluster availability as any node can process any query. Besides, it provides a good basis for dynamic load balancing through workload redistribution, as any node has local access to the entire database. But in the case of a shared-nothing cluster, the need for full replication yields poor disk utilization.

In [38], we proposed a solution which combines physical and virtual partitioning to define table subsets. This solution yields flexibility in intra-query parallelism while optimizing disk space usage and data availability. To validate our solution, we implemented a Java prototype on the 32-node cluster of the Paris team at IRISA. Experiments with our partitioning technique using the TPC-H benchmark queries gave linear and super-linear speedup, thereby reducing significantly the time of typical OLAP heavy-weight queries.

6.4.4. Data management in APPA

Participants: Reza Akbarinia, Vidal Martins, Esther Pacitti, Patrick Valduriez.

APPA (Atlas Peer-to-Peer Architecture) is a new P2P data management system which we are building. Its main objectives are scalability, availability and performance for advanced applications. These applications must deal with semantically rich data (e.g., XML documents, relational tables, etc.) using a high-level SQL-like query language. As a potential example of advanced application that can benefit from APPA, consider the cooperation of scientists who are willing to share their private data (and programs) for the duration of a given experiment. APPA has a network-independent architecture in terms of basic and advanced services that can be implemented over different P2P networks (unstructured, DHT, super-peer, etc.). This allows us to exploit continuing progress in such systems. To deal with semantically rich data, APPA supports decentralised schema management, data replication and updates, query processing and load balancing.

An important aspect of collaborative applications is that users may perform updates and may join and leave the network whenever they wish, thus hurting data availability. Data replication can then be used to increase availability. Lazy master replication is not applicable in P2P because a single master peer hurts availability. Thus, we focus on a multi-master approach in which all peers may update replicated data. In particular, we are interested in optimistic solutions based on semantic reconciliation, because it provides more flexibility and supports connections and disconnections. Existing semantic reconciliation solutions are typically performed at a single peer which may become a bottleneck in a large-scale system. In [49], we proposed a Distributed Semantic Reconciliation Algorithm (DSR). DSR enables optimistic multi-master replication, and assures eventual consistency among replicas. We proved the algorithm's correctness and validated it through implementation and simulation. The performance results show that DSR outperforms centralized solutions by a factor of 1.5 and has good scale up.

High-level queries over a large-scale P2P system may produce very large numbers of results that may overwhelm the users. To avoid such overwhelming, a solution is to use top-k queries whereby the user can specify a limited number (k) of the most relevant answers. However, relying on centralized histograms no longer applies in the P2P case. In [64], we proposed a fully distributed algorithm for executing top-k queries in P2P systems. We presented our algorithm for the case of unstructured systems, thus with minimal assumptions regarding the network. Our algorithm requires no global information, does not depend on the existence of certain peers, its bandwidth cost is low and addresses the volatility of peers during query execution. We validated our algorithm through implementation over a 64-node cluster and simulation using the BRITE topology generator and SimJava. Our performance evaluation shows that our algorithm has logarithmic scale up and improves Top-k query response time very well using P2P parallelism in comparison with baseline algorithms

7. Contracts and Grants with Industry

7.1. Microsoft Research (2003-2006)

Participants: Jean Bézivin, Patrick Valduriez.

The objective is to contribute to the development of the AMMA model management framework and foster the dissemination of our results as Open Source Software under a non restrictive license. In particular, we are adapting the AMMA framework to the principles and tools of the Microsoft Software Factory approach (Visual Studio 2005 Team System). Artefacts built by tools as ATL, AM3, AMW should be made available to the Microsoft environment with the help of technical space projectors.

7.2. IBM/Eclipse (2004-2005)

Participant: Jean Bézivin.

The objective is to port the ATL platform in the Eclipse Open Source environment. This was the only French project granted by IBM Eclipse Educational Grant in 2004. A first version of the prototype has been presented at the OOPSLA conference in october 2004 in Vancouver.

7.3. RNTL Modathèque (2004-2005)

Participants: Jean Bézivin, Patrick Valduriez.

In this project, we work with Thales RT (project leader), France Telecom R&D, LIP6 and software engineering tool vendors in France. The objective is to define the conceptual and practical basis for model engineering, in particular, using components of the Model Driven Environment (MDE) of the OMG. In this project, we use our ATL platform for MDE components.

7.4. Caroll Motor (2003-2006)

Participants: Jean Bézivin, Patrick Valduriez.

In the context of the Caroll joint venture between INRIA, CEA and Thales, the objective of the Motor project was to study the interoperability of model transformation languages. In this project, we showed interoperability results based on ATL. More generally, the principles of the AMMA platform are also being studied in this project.

7.5. IP Modelware (2004-2006)

Participants: Freddy Allilaire, Jean Bézivin, Peter Rosenthal, Patrick Valduriez.

In this very large european project, we work with Thales (project leader), IBM UK, IBM Israel, France Telecom R&D, LIP6 and the major industrial actors in model engineering in Europe. The objective is to demonstrate within 4 years the industrial application of model engineering.

7.6. Ilog (2005)

Participant: Jean Bézivin.

Ilog has hired one of our MS students for 3 months to study at LINA the applicability of model engineering and ATL to their tool suite.

8. Other Grants and Activities

8.1. Regional Actions

We are involved in two projects:

8.1.1. COM, Région Pays-de-la-Loire (2000-2006)

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.

8.1.2. GeoPict, Mégalis, Région Pays-de-la-Loire (2005-2006)

Participants: Marc Gelgon, José Martinez, Noureddine Mouaddib.

GeoPict is a joint project between Magic Instinct Software, a startup in Nantes, and three research teams at LINA and IRCCyN. The motivation is to take advantage of high-speed networks in order to create new services to transmit huge amounts of information such as multimedia data. The goal of the project is to provide an on-line service to access and visualise geo-referenced videos connected to a geographic information system (namely, the register – *cadastre* – of each city). More precisely, videos are recorded at 360 degrees while driving along the streets of a city. This information has to be stored, connected to the geographic database thanks to the spatial positioning recorded during the travelling. Next, the video information must be mixed with 3D geographical models in order to reconstruct panoramic views at any point in the city, as well as virtual reality trips.

8.2. National Actions

We are involved in three projects:

8.2.1. ACI MDP2P (2003-2006)

Participants: Gaëtan Gaumer, Esther Pacitti, Patrick Valduriez.

The project MDP2P (Massive data management in peer-to-peer systems) of the ACI Masses of Data of the French ministry of research is led by the Atlas group and involves three other INRIA groups: Paris and Texmex in Rennes, and Gemo in Orsay. The main objective of the project is to provide high-level services for managing text and multimedia data in large-scale P2P systems. Similar to database management systems, these services

are not limited to file sharing (like current P2P systems) and need be high-level with query capabilities and transactional support (for data consistency). Furthermore, they must provide good access performance which can be obtained through data replication, distributed query optimization, and parallel query processing. To validate our approach and show its wide range of application, we concentrate on two different P2P contexts that we know well: the Web and clusters of PC.

8.2.2. *ACI SemWeb (2004-2007)*

Participants: José Martinez, Noureddine Mouaddib, Guillaume Raschia.

The project SemWeb (Querying the Semantic Web with XQuery) of the ACI Masses of Data involves PRiSM, Versailles, CNAM, Paris, LIP6, Paris, SIS, Toulon and LINA, Nantes. The project aims at studying problems and providing solutions to XML-based mediators in the context of the Semantic Web using XQuery as the common querying language. Foreseen main problems are scalability of the proposed architecture, integration of heterogeneous sources of information, and dealing with metadata. The results of the project should be an homogeneous mediator architecture, exemplified on typical applications, and delivered as a open-source software.

8.2.3. *ACI APMD (2004-2007)*

Participants: José Martinez, Noureddine Mouaddib, Guillaume Raschia.

The project APMD (Personalised Access to Masses of Data)(2004-2007) of the ACI Masses of Data involves PRiSM, Versailles, CLIPS-IMAG, Grenoble, IRISA, Lannion, IRIT, Toulouse, LINA, Nantes and LIRIS, Lyon. The goal of the project is to improve the quality of retrieved information thanks to personalisation techniques or, in other words, to personalise the retrieved information in order to improve its quality with respect to the end user. This is of major importance for applications targeted to a large audience, like e-commerce, which have to take into account a large number of parameters: heterogeneous sources of information, various data formats, used languages, large amount of available data, etc. More precisely, the project has to define precisely which are the components of a user's profile, how it can evolve, and then take advantage of these profiles in order to filter and present adaptively the retrieved information, especially when dealing with huge amounts of information.

8.3. International actions

We are involved in the following international actions:

- the Interop European network of excellence (2003-2006) with all the research groups working on model engineering in Europe;
- the Daad (Distributed computing with Autonomous Applications and Databases) project (2003-2007), funded by CAPES in Brazil and COFECUB in France, with UFRJ, Brazil, on distributed data management;
- the GridData project (2005-2008) Databases) project, funded by CNPQ in Brazil and INRIA, with the Gemo project-team and the universities PUC-Rio and UFRJ, Brazil, on data management in Grid environments;
- 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 William 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 organising the French database research community, in the context of the I3 GDR and of the conference Bases de Données Avancées (BDA). P. Valduriez is a member of the scientific committee of the ACI GRID and a member of the ACM SIGMOD steering committee.

J. Bézivin is a member and co-founder of the steering committee of the ECOOP (AITO) and UML/Models conferences. In 2006, he is a co-organizer of a track on model transformation of the ACM Symposium of Applied Computing to be held in Dijon.

The 20th ECOOP conference will take place in Nantes in July 2006. The conference co-chairs are J. Bézivin and Pierre Cointe (Obasco project team in Nantes) who founded ECOOP 20 years ago. In addition to their international recognition, this also demonstrates the strong cooperation between Atlas and Obasco. ECOOP is currently ranked 39th on a total of more than 1200 conferences and computer science journals by CiteSeer.

In 2008, the Atlas group will organize the EDBT conference in Nantes. N. Mouaddib and P. Valduriez have accepted to be executive chairs and they will present a proposal to the steering committee at EDBT 2006. EDBT is currently ranked 170th (top 13 percent) on Citeseer.

9.2. Editorial Program committees

Participation in the editorial board of scientific journals:

- Distributed and Parallel Database Systems, Kluwer Academic Publishers: P. Valduriez
- Internet and Databases: Web Information Systems, Kluwer Academic Publishers: P. Valduriez
- Ingenierie des Systèmes d'Information, Hermès : N. Mouaddib, P. Valduriez
- Journal of Object Technology: J. Bézivin
- SoSyM, Software and System Modeling, Springer Verlag: J. Bézivin
- IEEE Transactions Journal on Fuzzy Systems : N. Mouaddib
- International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems: N. Mouaddib

Participation in conference programme committees :

- Int. Conf. on Very Large Databases (VLDB), 2005, 2006: P. Valduriez
- Int. Conf. on Parallel and Distributed Computing (Euro-Par), 2005: P. Valduriez (vice chair, Distributed and Parallel Databases track)
- Int. Conf. on High Performance Computing for Computational Science (VecPar), 2006: P. Valduriez
- Int. Workshop on High-Performance Data Management in Grid Environments (HPDGrid 2006), co-located with VecPar 2006: P. Valduriez (co-chair with M. Mattoso, UFRJ), E. Pacitti (PC chair), G. Raschia (PC member)

- Journées Bases de Données Avancées (BDA), 2005: N. Mouaddib, E. Pacitti
- Brazilian Symposium on Databases (SBBD), 2005: E. Pacitti
- Journées Francophones sur la Cohérence des Données en Univers Réparti (CDUR), IEEE France et ACM SIGOPS France: E. Pacitti
- Conférence sur la Recherche d'Information et Applications (CORIA), 2005: J. Martinez
- ACM Int. Symp. on Applied Computing: Multimedia and Visualisation Track (SAC), 2004: J. Martinez
- Int. Conf. on Enterprise Information Systems (ICEIS), 2005, 2006: J. Bézivin
- Enterprise Distributed Object Computing (EDOC), 2005, 2006: J. Bézivin
- Fundamental Approaches to Software Engineering (ETAPS/FASE), 2005, 2006: J. Bézivin

9.3. Invited Talks

In July, E. Pacitti and P. Valduriez visited the University of Waterloo to work with Tamer Özsu and his team. E. Pacitti gave a lecture on preventive replication in cluster systems and P. Valduriez gave a lecture on top-K query processing in P2P systems, both in the database seminar series.

E. Pacitti gave a talk on preventive replication and the RepDB* prototype at UFRJ, Rio de Janeiro.

P. Valduriez gave an invited talk on large-scale experimentation with preventive replication at the workshop on Design, Deployment and Implementation of Database Replication at VLDB 2005.

J. Bézivin gave an invited talk at Fujaba Days, Paderborn, Germany, and another at Object Days, Erfurt, Germany.

9.4. Teaching

All the members of the Atlas group teach database management, multimedia, and softwareengineering 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] 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 VLDB 2000", vol. 10, n° 2-3, 2001.
- [2] T. ÖZSU, P. VALDURIEZ. *Principles of Distributed Database Systems, 2nd edition*, Prentice Hall, 1999.

Books and Monographs

- [3] J. BÉZIVIN, M. DIDONET DEL FABRO, F. JOUAULT, I. KURTEV, P. VALDURIEZ. *Model Engineering: From Principles to Applications*, to appear, Springer Verlag, 2006.

Doctoral dissertations and Habilitation theses

- [4] N. DESSAIGNE. *Le modèle DOAN (DOcument Annotation Model): modélisation de l'information complexe appliquée à la plateforme Arisem Kaliwatch Server*, Ph. D. Thesis, Adv. J. Martinez, University of Nantes, France, dec 2005.
- [5] F. EVEN. *Extraction d'information et modélisation de connaissances à partir de notes de communication orale*, Ph. D. Thesis, Adv. N. Mouaddib, University of Nantes, France, oct 2005.
- [6] C. LEPAPE. *Contrôle de qualité des données répliquées dans un cluster*, Ph. D. Thesis, Adv. P. Valduriez and S. Gançarski, University Paris 6, France, dec 2005.
- [7] E. LOISANT. *Navigation dans une classification de collection d'images*, Ph. D. Thesis, Adv. J. Martinez, University of Nantes, France, dec 2005.
- [8] A. PIGEAU. *Structuration geo-temporelle de données multimédia personnelles pour la navigation sur terminal mobile*, Ph. D. Thesis, Adv. N. Mouaddib and M. Gelgon, University of Nantes, France, dec 2005.
- [9] R. SAINT-PAUL. *Une architecture pour le résumé en ligne de données relationnelles et ses applications*, Ph. D. Thesis, Adv. N. Mouaddib and G. Raschia, University of Nantes, France, jul 2005.

Articles in refereed journals and book chapters

- [10] F. ALLILAIRE, J. BÉZIVIN, M. DIDONET DEL FABRO, F. JOUAULT, D. TOUZET, P. VALDURIEZ. *AMMA : vers une Plate-forme Générique d'Ingénierie des Modèles*, in "Génie Logiciel", n° 73, 2005, p. 8-15.
- [11] B. BENATALLAH, M. HASSAN, N. MOUADDIB, H. PAIK, F. TOUMANI. *Building and Querying an e-Catalog Network Using P2P and Data Summarisation Techniques*, in "Journal of Intelligent Information Systems (JIIS)", to appear, 2005.
- [12] J. BÉZIVIN. *On the Unification Power of Models*, in "Software and Systems Modelling (SoSym)", vol. 4, n° 2, 2005, p. 171-188.
- [13] S. GANÇARSKI, H. NAACKE, E. PACITTI, P. VALDURIEZ. *The Leganet System. Freshness-Aware Transaction Routing in a Database Cluster*, in "Information Systems", to appear, 2005.
- [14] M. GELGON, P. BOUTHEMY, J.-P. LE CADRE. *Recovering and Associating the Trajectories of Multiple Moving Objects in an Image Sequence with a PMHT Approach*, in "Image and Vision Computing", vol. 23, n° 1, 2005, p. 19-31.
- [15] S. HAMMOUDI, D. LOPES, J. BÉZIVIN. *Approche MDA pour le Développement d'Applications Internet sur des Plates-formes de Service web: modélisation, transformation et prototypage*, in "Ingénierie des Systèmes d'Information", vol. 10, n° 3, 2005, p. 67-90.
- [16] C. LEPAPE, S. GANÇARSKI, P. VALDURIEZ. *Data Quality in a Database Cluster with Lazy Replication*, in "Journal of Distributed Information Management", vol. 3, n° 2, 2005.

- [17] E. LOISANT, J. MARTINEZ, H. ISHIKAWA, K. KATAYAMA. *Galois Lattices as a Classification Technique for Image Retrieval*, in "IPSJ Transactions on Data", n° 28, 2005.
- [18] P.-A. MULLER, P. STUDER, F. FONDEMENT, J. BÉZIVIN. *Platform independent Web Application Modeling and Development with Netsilon*, in "Software and Systems Modelling (SoSym)", to appear, 2005.
- [19] E. PACITTI, C. COULON, P. VALDURIEZ, T. ÖZSU. *Preventive Replication in a Database Cluster*, in "Distributed and Parallel Databases", to appear, 2005.
- [20] A. VOGLOZIN, G. RASCHIA, L. UGHETTO, M. MOUADDIB. *Querying the SAINTETIQ Summaries: a first attempt*, in "Journal of Intelligent Information Systems (JIIS)", to appear, 2005.

Publications in Conferences and Workshops

- [21] A. ABOUZAHRA, J. BÉZIVIN, M. DIDONET DEL FABRO, F. JOUAULT. *A Practical Approach to Bridging Domain Specific Languages with UML profiles*, in "Workshop on Best Practices for Model Driven Software Development, OOPSLA 2005".
- [22] J. BÉZIVIN, A. BELANGOUR, M. FREDJ. *The Design Decision Assistant Tool: A Contribution to the MDA Approach*, in "Journée sur l'Ingénierie Dirigée par les Modèles (IDM 2005), Paris", 2005.
- [23] J. BÉZIVIN, H. BRUNELIÈRE, F. JOUAULT, I. KURTEV. *Model Engineering Support for Tool Interoperability*, in "Workshop in Software Model Engineering (WiSME 2005), Montego Bay, Jamaica", 2005.
- [24] J. BÉZIVIN, C. BRUNETTE, R. CHEVREL, F. JOUAULT, I. KURTEV. *Bridging the Generic Modeling Environment (GME) and the Eclipse Modeling Framework (EMF)*, in "Workshop on Best Practices for Model Driven Software Development, OOPSLA 2005".
- [25] J. BÉZIVIN. *Some Lessons Learnt in the Building of a Model Engineering Platform*, in "Workshop in Software Model Engineering (WiSME 2005)".
- [26] J. BÉZIVIN. *Model Driven Engineering: Principles, Scope, Deployment and Applicability*, in "Generative and Transformational Techniques in Software Engineering (GTTSE'05), Tutorials, Braga, Portugal", 2005, p. 1-33.
- [27] J. BÉZIVIN, V. DEVEDZIC, D. DJURIC, J. FAVREAU, D. GASEVIC, F. JOUAULT. *An M3-Neutral Infrastructure for Bridging Model Engineering and Ontology Engineering*, in "1st Int. Conf. on Interoperability of Enterprise Software and Applications (INTEROP-ESA 2005), Geneva, Switzerland", Springer, 2005.
- [28] J. BÉZIVIN, F. JOUAULT. *Using ATL for Checking Models*, in "Int. Workshop on Graph and Model Transformation (GraMoT), Tallinn, Estonia", 2005.
- [29] J. BÉZIVIN, F. JOUAULT, P. ROSENTHAL, P. VALDURIEZ. *Modeling in the Large and Modeling in the Small*, in "European MDA Workshops: Foundations and Applications, MDFAFA 2003 and MDFAFA 2004", vol. LNCS 3599, Springer, 2005, p. 33-46.

-
- [30] J. BÉZIVIN, F. JOUAULT, D. TOUZET. *Principles, Standards and Tools for Model Engineering*, in "Workshop on Using Metamodels to Support MDD, IEEE Int. Conf. on Engineering of Complex Computer Systems (ICECCS 2005), Shanghai, China", 2005.
- [31] C. COULON, E. PACITTI, P. VALDURIEZ. *Consistency Management for Partial Replication in a High Performance Database Cluster*, in "IEEE Int. Conf. on Parallel and Distributed Systems (ICPADS 2005), Fukuoka, Japan", 2005.
- [32] C. COULON, E. PACITTI, P. VALDURIEZ. *Optimistic Preventive Replication in a Database Cluster*, in "Journées Bases de Données Avancées (BDA05), Saint Malo", 2005.
- [33] C. COULON, E. PACITTI, P. VALDURIEZ. *Scaling up the Preventive Replication of Autonomous Databases in Cluster Systems*, in "Int. Conf. on High Performance Computing for Computational Science (VecPar 2004)", vol. LNCS 3402, Springer, 2005, p. 174-188.
- [34] N. DESSAIGNE, J. MARTINEZ. *A Model for Describing and Annotating Documents*, in "European-Japanese Conference on Information Modelling and Knowledge Bases (EJC 2005), Tallinn, Estonia", 2005.
- [35] M. DIDONET DEL FABRO, J. BÉZIVIN, F. JOUAULT, E. BRETON, G. GUELTAS. *AMW: a generic model weaver*, in "Journée sur l'Ingénierie Dirigée par les Modèles (IDM 2005), Paris", 2005.
- [36] M. DIDONET DEL FABRO, J. BÉZIVIN, F. JOUAULT, P. VALDURIEZ. *Applying Generic Model Management to Data Mapping*, in "Journées Bases de Données Avancées (BDA05), Saint Malo", October 2005.
- [37] M. DIDONET DEL FABRO, F. JOUAULT. *Model Transformation and Weaving in the AMMA Platform*, in "Generative and Transformational Techniques in Software Engineering (GTTSE'05), Workshop, Braga, Portugal", July 2005, p. 71-77.
- [38] C. FURTADO, A. LIMA, E. PACITTI, P. VALDURIEZ, M. MATTOSO. *Physical and Virtual Partitioning in OLAP Database Clusters*, in "Int. Symp. on Computer Architecture and High Performance Computing, Rio de Janeiro, Brazil", 2005.
- [39] N. IDRISSE, J. MARTINEZ, D. ABOUTAJDINE. *On the Use of Co-occurrence Matrices for Texture Description in Image Databases*, in "Int. Conf. on Modeling and Simulation (ICMS 2005), Marrakech, Morocco", 2005.
- [40] N. IDRISSE, J. MARTINEZ, D. ABOUTAJDINE. *Selecting a Discriminant Subset of Co-occurrence Matrix Features for Texture-based Image Retrieval*, in "Int. Symp. on Visual Computing ISVC'05, Lake Tahoe, Nevada, USA", 2005.
- [41] N. IDRISSE, J. MARTINEZ, D. ABOUTAJDINE. *Selection of a Discriminant Subset of Co-occurrence Matrix Features for Image Search-by-Content*, in "IEEE Int. Conf. on Computer Systems and Information Technology (ICSIT 2005), Algiers, Algeria", 2005.
- [42] F. JOUAULT. *Loosely Coupled Traceability for ATL*, in "Workshop on Traceability, European Conference on Model Driven Architecture (ECMDA 2005), Nuremberg, Germany", 2005.

-
- [43] F. JOUAULT, I. KURTEV. *Transforming Models with ATL*, in "Workshop on Model Transformations in Practice, MoDELS 2005, Montego Bay, Jamaica", 2005.
- [44] F. JOUAULT, I. KURTEV. *On the Architectural Alignment of ATL and QVT*, in "Proceedings of ACM Symposium on Applied Computing (SAC 06), model transformation track, Dijon, France", To appear, April 2006.
- [45] I. KURTEV, K. VAN DEN BERG, F. JOUAULT. *Rule-based Modularization in Model Transformation Languages illustrated with ATL*, in "Proceedings of ACM Symposium on Applied Computing (SAC 06), model transformation track, Dijon, France", To appear, April 2006.
- [46] C. LEPAPE, S. GANÇARSKI, P. VALDURIEZ. *Replica Refresh Strategies in a Database Cluster*, in "Journées Bases de Données Avancées (BDA05), Saint Malo", 2005.
- [47] E. LOISANT, J. MARTINEZ, H. ISHIKAWA, M. OHTA, K. KATAYAMA. *Building an Image Navigation Structure from a Fuzzy Relationship*, in "IEEE Int. Workshop on Managing Data for Emerging Multimedia Applications (EMMA 2005), Tokyo, Japan", 2005.
- [48] D. LOPES, S. HAMMOUDI, J. BÉZIVIN, F. JOUAULT. *Generating Transformation Definition from Mapping Specification: Application to Web Service Platform*, in "Int. Conf. on Advanced Information Systems Engineering (CAiSE 2005), Porto, Portugal", 2005.
- [49] V. MARTINS, E. PACITTI, P. VALDURIEZ. *Distributed Semantic Reconciliation of Replicated Data*, in "Journées Francophones sur la Cohérence des Données en Univers Réparti (CDUR), IEEE France et ACM SIGOPS France, Paris", 2005.
- [50] I. MBAYE, O. HAJ THAMI, J. MARTINEZ. *A Model for Indexing Videos and Still Images from the Moroccan Cultural Heritage*, in "IEEE Int. Workshop on Multimedia Signal Processing MMSP 2005, Shanghai, China", 2005.
- [51] L. NAOUM, G. RASCHIA, N. MOUADDIB. *Manipulating Fuzzy Summaries of Databases: Unary Operators and their Properties*, in "Joint EUSFLAT & LFA Conference (EUSFLAT-LFA 2005), Barcelona, Spain", 2005.
- [52] A. NIKSERESHT, M. GELGON. *Agrégation légère de mélange de lois gaussiennes en vue de l'indexation multimédia répartie*, in "Journées France Telecom CORESA, Rennes", 2005.
- [53] A. NIKSERESHT, M. GELGON. *Low-cost Distributed Learning of a Gaussian Mixture Model for Multimedia Content-based Indexing on a P2P Network*, in "ACM Workshop on Multimedia Information Retrieval (MIR 2005), Singapore", 2005.
- [54] A. NIKSERESHT, M. GELGON. *Decentralized Distributed Learning of a Multimedia Class for Content-based Indexing*, in "Euromicro Conf. on Parallel, Distributed and Network-based processing (PDP 2006), Montbeliard", to appear, 2006.
- [55] A. PIGEAU, M. GELGON. *Building and Tracking Hierarchical Partitions of Image Collections on Mobile Devices*, in "ACM Multimedia conference, Singapore", 2005.

- [56] J. ROUGUI, M. GELGON, M. RZIZA, J. MARTINEZ, D. ABOUTAJDINE. *Hierarchical organization of a set of Gaussian mixture speaker models for scaling up indexing and retrieval in audio documents*, in "Proc. of ACM Symp. on Applied Computing (SAC'2006), Dijon, France", April 2006.
- [57] R. SAINT-PAUL, G. RASCHIA, N. MOUADDIB. *Database Summarization: The SaintEtiQ System*, in "Journées Bases de Données Avancées (BDA05), Saint Malo", demo, 2005.
- [58] R. SAINT-PAUL, G. RASCHIA, N. MOUADDIB. *General Purpose Database Summarization*, in "Int. Conf. on Very Large Databases (VLDB 2005), Trondheim, Norway", Morgan Kaufmann Publishers, 2005, p. 733–744.
- [59] R. SAINT-PAUL, G. RASCHIA, N. MOUADDIB. *Résumé généraliste de bases de données*, in "Journées Bases de Données Avancées (BDA05), Saint Malo", 2005.
- [60] P. VALDURIEZ, E. PACITTI, C. COULON. *Large-scale Experimentation with Preventive Replication in a Database Cluster*, in "Workshop on Design, Deployment and Implementation of Database Replication, VLDB 2005, Trondheim, Norway", 2005.
- [61] P. VALDURIEZ, E. PACITTI. *Data Management in Large-scale P2P Systems*, in "Int. Conf. on High Performance Computing for Computational Science (VecPar 2004)", vol. LNCS 3402, Springer, 2005, p. 109-122.
- [62] A. VOGLOZIN, G. RASCHIA, L. UGHETTO, M. MOUADDIB. *Querying the SaintEtiQ Summaries – Dealing with Null Answers*, in "IEEE Int. Conf. on Fuzzy Systems (FUZZ-IEEE 2005), Reno, Nevada, USA", 2005.

Miscellaneous

- [63] R. AKBARINIA, V. MARTINS, E. PACITTI, P. VALDURIEZ. *Data Management in the APPA P2P System*, 2005, <http://www.sciences.univ-nantes.fr/lina/ATLAS/MDP2P/>, Journées Masses de Données en Pair-à-Pair.
- [64] R. AKBARINIA, E. PACITTI, P. VALDURIEZ. *Top-n Query Processing in the APPA P2P System*, 2005, <http://www.sciences.univ-nantes.fr/lina/ATLAS/MDP2P/>, Journées Masses de Données en Pair-à-Pair.
- [65] R. HAYEK. *Vers un service de résumé avancé dans APPA*, sep 2005.
- [66] J. MARTINEZ. *Gestion de bases d'images sur un réseau : un panorama de propositions*, 2005, <http://www.sciences.univ-nantes.fr/lina/ATLAS/MDP2P/>, Journées Masses de Données en Pair-à-Pair.
- [67] V. MARTINS, E. PACITTI, P. VALDURIEZ. *Replication in the APPA P2P System*, 2005, <http://www.sciences.univ-nantes.fr/lina/ATLAS/MDP2P/>, Journées Masses de Données en Pair-à-Pair.