INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team ATLAS

Complex Data Management in Distributed Systems

Rennes - Bretagne Atlantique

THEME SYM

Table of contents

1. Team

Head of Project Team
Patrick Valduriez [Director of Research, HdR]

Faculty Members, University of Nantes
Jean Bézivin [Professor, HdR]
Marc Gelgon [Associate Professor, HdR]
José Martinez [Professor, HdR]
Noureddine Mouaddib [Professor, HdR]
Esther Pacitti [Associate Professor]
Guillaume Raschia [Associate Professor]

Technical Staff
Freddy Allilaire [Engineer, Modelplex]
Hugo Brunelière [Engineer, Modelplex]
William Kokou Dedzoe [Engineer, Grid4All since dec.]
Brahim-Khalil Loukil [Engineer, ODL since oct.]
Mounir Tlili [Engineer, Xwiki Concerto since feb.]
Sarfraz Ashfaq [Engineer, Grid4All from feb. to nov.]

Administrative Assistant
Élodie Lizé [Secretary, CDD]

Ph.D. Students
Reza Akbarinia [Fellowship INRIA]
Mikaël Barbero [Fellowship Modelplex]
Marcos Didonet Del Fabro [Fellowship INRIA until sep.]
Rabab Hayek [Fellowship MENRT]
Manal El Dick [Fellowship MENRT]
David Faye [Fellowship Egide]
Sandra Lemp [ATER U. Nantes until sep.]
Vidal Martins [Fellowship PUCPR until may]
Jorge Manjarrez Sanchez [Fellowship Conacyt]
Afshin Nikseresht [Fellowship Iranian gov.]
Jorge Quiane Ruiz [Fellowship Conacyt]
Jamal Rougui [Egide Fellowship]
Anthony Ventresque [Fellowship MENRT]
Amenel Voglozin [ATER U. Nantes]

Post-doctoral fellows
Marcos Didonet Del Fabro [Fellowship INRIA/ILOG since oct.]
Frédéric Jouault [Fellowship OpenEmbeDD, U. Alabama at Birmingham (USA)]
Hanna Kozankiewicz [Fellowship ERCIM until aug.]
Kwangjin Park [Fellowship INRIA until aug.]

2. Overall Objectives

2.1. Introduction

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 management difficult. Third, this combined evolution of data and systems gives rise to new, typically complex, applications with ubiquitous, on-line data access: collaborative content management (e.g. Wiki), virtual libraries, virtual stores, global catalogs, services for personal content management, etc.

The general problem can be summarized as complex data management in distributed systems. The Atlas project-team 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 themes. The theme "database summaries" addresses the issues of data abstraction from large, distributed 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 inter-theme research activities are: multimedia data management in cluster systems, database summaries in P2P systems, and model management applied to distributed data integration.

2.2. Highlights of the year

This year has been very productive in terms of research results and technology transfer, with:

- a very good publication record, including 3 papers published in the top two database conferences (1 SIGMOD, 2 VLDB); a strong result is the best position algorithms for top-k queries which significantly outperform the best algorithms published so far [33];
- 1 HDR and 6 Ph.D. theses defended;
- ATL, our transformation-based model management framework has been recognized as a standard component in Eclipse and is now integrated into the new M2M (Model-to-Model) project.

3. Scientific Foundations

3.1. Data Management

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

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

The fundamental principle behind data management is data abstraction, which enables applications and users to deal with the data at a high conceptual level while ignoring implementation details. The relational model, by resting on a strong theory (set theory and first-order logic) to provide data independence, has revolutionized data 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 rules and triggers, transaction support, etc.

This data independence principle has also enabled DBMS to continuously integrate new advanced capabilities such as object 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 has 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 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.

3.2. Multimedia Data Management

Multimedia data such as image, audio or video differs much from structured data and semi-structured (text) data in that they are 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: multimedia data analysis and pattern recognition, information retrieval and distributed databases. Multimedia data management is organized around the three fundamental parts of database design: modelling, querying and indexing. However, they have to be considerably adapted in order to maintain 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 transformations on the original data (e.g. image texture features), that translate into feature vectors, to more abstract representations (e.g. parametric models), that attempt to capture a class rather than an instance of multimedia elements. Furthermore, media content creators may add metadata information that conveys more semantics.Thus, multimedia data analysis deals with the design of suitable observations from multimedia and pattern recognition techniques.

Querying has been concerned with access by a high-level (SQL-like) query language on user-defined schemas. In contrast, techniques for querying multimedia data come from the information retrieval community. Athough extensible, each content-based multimedia system relies on a single, well-defined schema. 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. Thus, multimedia querying must be an interactive activity.

Indexing deals 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. Classical indexing structures (tree-based and hashing-based) are not effective, because of the "dimensionality curse problem", which states that the performance of indexing (and thus querying) degrades *severely* as the data dimensionality increases. 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.

3.3. Fuzzy Logic

The ever growing size of databases has made data summarization attractive in order to present the user a concise and complete view of the database. Our proposed summarization process [10] can roughly be described as a two step process. The first step is to rewrite the original database records into a unified user-oriented vocabulary. The second step is to use a concept formation algorithm against the rewritten data. The fuzzy set theory provides mathematical foundations to manage these two steps in a more user-friendly and robust way than the one that can be achieved with first order logic. Fuzzy sets theory was introduced by L.A. Zadeh in 1965 to model sets whose boundaries are not sharp. A fuzzy (sub)set F of a universe Ω is defined using 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)).

During summarization, 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 express their knowledge. A database user usually refers to her data using a vocabulary appropriate for her field of expertise and understood by her fellows. For example, a salary will be said to be high, reasonable or average. This description is an implicit categorization and there is no strict border line between an average and a high salary. Fuzzy logic offers the mathematical ground to define such a vocabulary in terms 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 the hierarchy should 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, a concept formation algorithm is less prone to suffer the well known threshold effect of similar incremental algorithms.

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 flexible querying and approximate answering. 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 allows handling user queries expressed in vague or imprecise terms.

3.4. Model Engineering

A model is a formal description of a design artefact 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 (DSL) that help capturing the various aspects of complex systems. Models are 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. 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 a 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.

3.5. Distributed Data Management

The Atlas project-team considers data management in the context of distributed systems, with the objective of making distribution transparent to the users and applications. Thus, we capitalize on the principles of distributed systems, in particular, large-scale distributed systems such as clusters, grid, and peer-to-peer (P2P) systems, to address issues in data replication and high availability, 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) [12]. 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. Overview

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 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 large-scale distributed collaborative applications, large decision-support applications or multimedia personal databases.

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. 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 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. Content-based retrieval from distributed multimedia documents is also an important class of applications.

5. Software

5.1. ATL (Atlas Transformation Language)

Participants: Jean Bézivin [contact], Frédéric Jouault, Patrick Valduriez.

URL: http://www.eclipse.org/m2m/atl/

ATL is a transformation-based model management framework, with metadata management and data mapping 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. It comes with a library of more than 100 transformation components. ATL is released as Open Source Software under the Eclipse Public Licence and available as an Eclipse plugin. There is now an active community of more than 100 user sites, including research labs and major companies (Airbus, NASA, Ilog, Sodius, Obeo, TNI, etc.). ATL has been selected by two competitivity centers in France: System@tic and TopCased. In 2007, ATL has been recognized a standard component in Eclipse and is now integrated into the new M2M (Model-to-Model) project.

5.2. AMW (Atlas Model Weaver)

Participants: Jean Bézivin [contact], 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 separate 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. AMW is released as Open Source Software under the Eclipse Public Licence and available as an Eclipse plugin. AMW is being used by more than 40 user sites, including research labs and major companies (NASA, BAE, Versata, Obeo, etc.).

5.3. APPA (Atlas Peer-to-Peer Architecture)

Participants: Reza Akbarinia, William Kokou Desdoe, Esther Pacitti, Mounir Tlili, Patrick Valduriez [contact].

URL: http://www.sciences.univ-nantes.fr/lina/gdd/appa/

APPA is a P2P data management system that provides scalability, availability and performance for applications which deal with semantically rich data (XML, relational, etc.). APPA provides advanced services such as queries, replication and load balancing. It is being implemented on top on various P2P networks such as JXTA and OpenChord and tested on GRID5000. Two new services have been implemented this year: KTS and SbQA. KTS (Key-based Timestamp Service) is a distributed service to manage timestamps in DHTs. It is useful to solve various DHT problems which need a total order on operations performed on each data, e.g. data currency. KTS is used in the Strep Grid4All and RNTL Xwiki Concerto projects as the basis to perform reconciliation of replicated documents in a P2P wiki system. SbQA is a Satisfaction-based Query Allocation framework for P2P environments where consumers and providers are autonomous and have special interests towards providers and queries, respectively. It is used in the Grid4All project as the basis for the resource discovery service.

5.4. DBSum

Participants: Noureddine Mouaddib, Guillaume Raschia [contact], Amenel Voglozin.

URL: http://www.lina.sciences.univ-nantes.fr/grim/doku.php?id=dbsum

DBSUM is a *Database Summary Management System* that provides various tools to support data reduction with query and analytical processing techniques on top of a DBMS. The current implementation has two parts: a summarization engine, namely SAINTETIQ, for building and updating database summaries; a full-feature user interface coined SEQT (*Summary Exploration and Querying Tool*) which provides languages, algorithms and views to query, search and browse into summaries. SAINTETIQ computes and maintains abstract and user-friendly views from very large databases. Aa an alternative to the win32 executable version of SAINTETIQ, SAINTETIQ is also exposed as a Web Service. SEQT is a new software component which provides efficient search algorithms to filter summaries and support flexible query processing and personalized queries.

6. New Results

6.1. Database summaries

Database summaries provide an effective way to access compact representations of large databases. The basis for our work is the DBSUM database summary management system. DBSUM capitalizes on our previous work on SAINTETIQ, a service for constructing fuzzy linguistic database summaries, as stated in Section 3.3. To deal with very large databases in DBSUM, we studied how our summary-based access method can satisfy the user's requirements by the way of flexible querying and deal with summaries in distributed systems.

6.1.1. Summary-based access method

Participants: Noureddine Mouaddib, Guillaume Raschia, Amenel Voglozin.

We completed our work on a flexible querying system that provides approximate answering capabilities using a summary-based access method . User's requirements are formulated into fuzzy linguistic predicates of conjunctive queries such as *Size is small* rather than regular range queries (*Size in* $[1, 5]$). And the system provides intentional answers to the end-user with respect to the descriptions of database summaries. The main underlying structure is a hierarchy of multidimensional summaries for which we have an efficient construction and maintenance algorithm. Experiments with this algorithm have shown excellent results [56] in terms of computational time, structural stability of the hierarchy and scaling up in the number of database records. We proposed a personnalized querying framework for the summaries and studied the coupling of our summary-based access method with the PostGreSQL database management system [19], [32]. This work leads to new and exciting research issues for database summaries in the areas of preference handling, ranking answers, and dealing with empty- and many-answer problems.

We started investigating the problem of setting the compression rate of a database summary from a given hierarchy. The main issue is to retrieve a rooted subtree in the overall summary hierarchy given a generalization position. This issue is difficult since every cut in the tree is a summary candidate, i.e. an exact cover of the database, from the root to the set of leaves. We studied the algebraic structure underlying the partial ordered set of cuts in a summary tree and proved important properties of the lattice of tree cuts. We extended SAINTETIQ's query engine such that it is possible to complete the answer of a given query and provide an exact cover of the data, i.e. a compressed table that meets the user's requirements formulated in the query predicates and is an abstract view of the whole database [59].

6.1.2. Summaries in distributed systems

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

In a distributed system with multiple data sources, we must deal with two problems: merging summaries from multiple sources and managing distributed summaries. The first problem can be defined as merging two summary hierarchies obtained from distinct data sources which have the same schema. We proposed various algorithms to generate a single summary hierarchy given two distinct hierarchies, without scanning the raw data [38]. We first studied the Greedy Merging Algorithm (GMA) that takes all the leaves of both hierarchies and generates the optimal partitioning for the union of two data sources. However, GMA has exponential time complexity. We then proposed two alternative solutions with constant time complexity w.r.t. the number of data items: the Merge by Incorporation Algorithm (MIA) which incorporates leaves of a tree inside the other one, as done by the SAINTETIQ engine, and the Merge by Alignment Algorithm (MAA) which consists in rearranging summaries by levels in a top-down manner.

We addressed the second problem in the context of P2P systems. In this context, summaries have two main virtues. First, they can be directly queried and used to approximately answer a query without exploring the original data. Second, as semantic indexes, they can help locating relevant nodes based on data content. We proposed to construct a global summary that describes all the data shared in the P2P network [44]. This global summary is stored in a distributed fashion. For a given query, the global summary can be used to determine the set of nodes having relevant data. Then, those targeted nodes are directly contacted. Simulation results have shown that the cost of query routing is significantly reduced compared to the flooding approach [45]. However, maintaining such a global summary is hard and costly in a P2P environment. To address this problem, in the context of unstructured networks, we proposed a connectivity-based partitioning of the network in peer domains, such that each domain is able to self-manage its global summary [45]

6.2. Model management

Developers of modern information systems must typically deal with different models and perform transformations between models, e.g. mapping heterogeneous data source descriptions in a global schema or converting XML documents into relational data. Today, most of these transformations are still programmed using specific languages like SQL, XSLT or even Java, Perl, or C. Model management provides a general approach to

this problem with techniques and tools for managing models and model transformations. To contribute to this approach, we have defined the Atlas Model Management Architecture (AMMA) which includes two major tools (ATL and AMW) released as Eclipse components. This year, the main results are on the management (including maintenance and specialization) of ATL transformations and on model integration with AMW.

6.2.1. Management of ATL transformations

Participants: Freddy Allilaire, Mikaël Barbero, Jean Bézivin, Frédéric Jouault.

ATL is a rule-based language for model transformation. Like the source and target models, an ATL transformation program is itself a model and thus conforms to a given metamodel. This provides much genericity and flexibility to manage ATL transformations. ATL tools now provide full support for the major tasks involved in using a language: editing, compiling, executing, and debugging [20].

Managing ATL transformations requires the ability of modularizing large transformations and extending them. In [25], we studied the problem of modularizing transformations in current rule-based model transformation languages such as ATL, with adaptability and reusability as main requirements. We experimented two scenarios by applying different transformation techniques: use of explicit and implicit rule calls, and use of rule inheritance. The experience with these scenarios shows that current languages provide a reasonably full set of modular constructs but may have problems in handling some composition tasks. In [37], we proposed a conceptual and practical approach to model extensibility, in which new models are created as derivations from base models.

Many research groups have been working on their own model transformation approaches and languages, like ATL. Thus, there is a growing need for transformation language interoperability. In [25], we proposed a set of heuristics to reason about the problems that must be addressed when translators between languages need to be developed. We also show that achieving a large degree of interoperability is difficult since some languages have incompatible features.

6.2.2. Model integration with AMW

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

Model transformations can be used in many different application scenarios, for instance, to integrate very large models coming from different databases. Thus, they are becoming more and more complex. However, there is not yet a complete solution that automates the development of model transformations. In [39], we proposed a novel approach that uses matching transformations and weaving models to semi-automate the development of transformations. Matching transformations are a special kind of transformations that implement heuristics to create weaving models, i.e., models that capture different kinds of relationships between models. Our solution enables to rapidly implement and customize these heuristics. The weaving models are derived into model integration transformations. To validate our approach, we presented a set of experiments with our ATL and AMW tools using metamodels with distinct size and complexity, which demonstrated the feasibility and scalability of our solution.

An important need for model integration is the capability to identify mappings and differences between models. There has been initial research toward model differentiation applied to UML diagrams, but differentiation of domain-specific models has not been explored deeply. In [27], we proposed metamodel-independent algorithms and associated tools for detecting mappings and differences between domain-specific models, with facilities for graphical visualization of the detected differences.

6.3. Multimedia data management

Recent advances such as joint text/audiovisual processing and semi-supervised learning are making large-scale indexing and retrieval of multimedia information a major research challenge. We address this challenge in the context of large-scale distributed systems. A first result is on scaling up multimedia data access by exploiting parallelism in a cluster system. A second result concerns estimation of probabilistic multimedia class models, by distributed learning.

6.3.1. Scaling up multimedia data access

Participants: Jorge Manjarrez, José Martinez, Patrick Valduriez.

We studied the general problem of content-based retrieval in the centralized approach, where a single database instance handles all users' queries. In the case of a large database where main-memory indexing is not possible, this approach is inherently limited by the cost of computing the similarity between the given query object and the database multimedia objects, as well as the disk access cost to retrieve the multimedia objects. We provided evidence of the performance limitations of this approach when a high-dimension index is used. As an alternative solution to indexing, we studied clustering which consists of partitioning the database objects into groups of similar objects on disk [48]. Above a number of dimensions (e.g. 30), clustering outperforms indexing. We analyzed the computational complexity of this approach and proposed a way to compute the optimal number of clusters which can lead to sublinear algorithms.

A natural extension of clustering to cope with very large databases is to use parallelism. However, a key aspect to efficiently exploit the performance offered by a parallel machine is data allocation. To this purpose, we defined a data allocation method for parallel database systems based on a shared-nothing (distributed) architecture [49]. This method enables efficient content-based retrieval of multimedia data and is general enough to be used in different parallel architectures.

6.3.2. Distributed learning of a probabilistic model for multimedia classification

Participants: Marc Gelgon, Afshin Nikseresht, Jamal Rougui.

Learning a probabilistic model that describes the distribution of multimedia class features, e.g. capturing visual variability, is a major task in multimedia information retrieval. Performing this task on distributed data sources opens new technical issues and applicative perspectives. Large vocabulary models require huge amounts of training data and are likely to be learned globally over the internet. We address the problem of aggregating local probabilistic models, in order to improve the estimated class-conditional distribution, which in turn decreases the probability of misclassification. We focus on distributions of the form of mixture models (mainly Gaussian). We defined a scheme for collective mixture learning from local mixture models, through mixture aggregation via model parameters. We propagate information in the network in an epidemic fashion.

Model aggregation relies on finding suitable combinations of incoming Gaussian components from multiple sources. To determine the number of components after aggregation, we proposed two solutions. The first one seeks a combination that preserves, as far as possible, the expected log-likelihood, averaged over incoming mixtures, using an approximation to the Kullback divergence between mixtures [29]. The second solution considers an Expectation-Maximization scheme operating on parameters of Gaussian components rather than on data. It provides a well-founded approach to estimate the combination and the number of components in the aggregated mixture model [43]. We also use a variation of the first solution for hierarchical partition tracking, in the context of personal image management [50], and for building an index structure over a set of mixture models [31]. Finally, we have started to explore an incremental, semi-supervised solution, for network surveillance applications [58].

6.4. Distributed data management

In a large scale distributed system, data sources are typically in high numbers, autonomous and very heterogeneous in size and complexity. Data management in this context offers new research opportunities since traditional distributed database techniques need to scale up while supporting data autonomy, heterogeneity, and dynamicity. We focus on peer-to-peer (P2P) systems with the development of the Atlas Peer-to-Peer Architecture (APPA) prototype. The main results this year are in the management of replicated data, top-k query processing and satisfaction-based query allocation.

6.4.1. Management of replicated data

Participants: Reza Akbarinia, Manal El Dick, Esther Pacitti, Patrick Valduriez.

We addressed three different problems related to the management of replicated data in P2P systems: improving the performance of reconciling divergent data; caching indexes of replicated files; and supporting data currency in Distributed Hash Tables (DHTs). In [40], we proposed a topology-aware approach that exploits physical topology information to perform P2P distributed data reconciliation, a major function for collaborative applications. Our algorithm relies on dynamically selecting nodes to execute specific steps while carefully placing relevant data. We showed that our algorithm introduces a gain of 50% compared to a base-line algorithm and still scales up.

In [41], we addressed the problem of reducing the P2P redundant traffic in file-sharing systems. We proposed a solution that performs index caching and efficient query routing while supporting keyword search. Our solution improves the probability of finding available copies of requested files by leveraging file replication. In addition, it can direct queries to close results, by using topological information in terms of file physical distribution.

In [34], [36], we addressed the problem of finding a current replica in DHTs. Supporting data currency in replicated DHTs is difficult as it requires the ability to return a current replica despite peers leaving the network or concurrent updates. Our solution relies on timestamping. For generating timestamps, we proposed a Key-based Timestamping Service (KTS) which performs distributed timestamp generation. Except for the cases where the availability of current replicas is very low, the expected number of retrieved replicas is typically small. Although not initially planned, we found that KTS can be useful to solve other DHT problems such as reconciliation of replicated data. And we now use it in two important project (Grid4All and Xwiki Concerto) to perform reconciliation of replicated documents in a P2P wiki system.

6.4.2. Top-k query processing

Participants: Reza Akbarinia, Esther Pacitti, Patrick Valduriez.

Top-k queries have received much interest in many different areas such as network and system monitoring, information retrieval and P2P systems because they avoid overwhelming the user with large numbers of uninteresting answers. The general problem of answering top-k queries can be modeled using lists of data items sorted by their local scores. The most efficient algorithm proposed so far for answering top-k queries over sorted lists is the Threshold Algorithm (TA). However, TA may still incur a lot of useless accesses to the lists. In [33], we proposed two new algorithms, BPA and BPA2, based on best positions, which stop much sooner. We proved that the execution cost of BPA can be (m-1) times lower than that of TA and that the cost of BPA2 can be about (m-1) times lower than that of BPA. Our performance evaluation, using several test databases, has confirmed our analysis.

In [35], we proposed an efficient mechanism for top-k query processing in DHTs. Although our algorithm is TA-style, it is much more general since it supports a large set of non monotone scoring functions including linear functions. We proved analytically the correctness of our algorithm and validated it through a combination of implementation and simulation. The results show very good performance, in terms of communication cost and response time.

6.4.3. Satisfaction-based query allocation

Participants: Jorge Quiane Ruiz, Patrick Valduriez.

In the context of dynamic distributed systems, with large numbers of heterogeneous, autonomous information sources, participants can play two roles: consumers that generate queries and providers which perform the queries to give answers. In this context, we considered the query allocation problem. A main objective in query allocation is to obtain good response time. Most of the work towards this objective has dealt with the problem of load balancing. However, participants may have some interests towards providers and queries that are not only performance-related.

In [53], [51], we proposed a query allocation method that balances queries among providers by considering the providers' interests in addition to their load. We compared our solution to a well-known query load balancing algorithm. The results show that our solution yields high efficiency while supporting the providers' interests. In [54], we proposed a model to characterize the participants' interests in the long-run. We proposed a query allocation framework that trades consumers' interests for providers' interests according to their satisfaction. We compared our framework proposal to both query load balancing and economic approaches. The experimentation results show that our approach significantly outperforms baseline approaches by guaranteeing good response times and interesting queries to providers even under low arrival query rates. In [52], we introduce the use of money in our query allocation framework proposal by explaining how providers compute their bids for queries and how our approach allocates queries by considering both consumers' interests and providers' bids.

7. Contracts and Grants with Industry

7.1. IP Modelplex (2006-2009)

Participants: Freddy Allilaire, Jean Bézivin, Mikael Barbero, Hugo Brunelière, Patrick Valduriez.

The project (with Thales, IBM, Sodifrance, SAP, etc.) aims at defining a coherent infrastructure for the development of complex systems, where complexity corresponds to several factors like size, heterogeneity, dynamic evolution, distribution and subsystem autonomy. Examples of highly heterogeneous systems are legacy systems that have been built and adapted on long period of time, using different technologies. Model driven reverse engineering is also an important problem which we address.

7.2. STREP Grid4All (2006-2008)

Participants: Reza Akbarinia, Sarfraz Ashfaq, William Kokou Dedzoe, Esther Pacitti, Jorge Quiane, Patrick Valduriez.

The project is with France Telecom R&D (leader), INRIA (Atlas, Grand-Large, Régal, and Sardes), Kungliga Tekniska Hoegskolan, Swedish Institute of Computer Science, ICCS (Greece), University of Piraeus Research Center, Universitat Politècnica de Catalunya and Rededia S.L. (Spain). Atlas and INRIA-Rennes are the INRIA representatives. The goal of Grid4All is to develop a grid infrastructure and middleware for the collaboration of dynamic, small virtual organizations such as communities, schools and families. The main technical innovation is to foster the combination of grid and P2P techniques to provide a light-weight, flexible solution. Atlas contributes to the definition of the P2P infrastructure (which is based on APPA) and to the development of two key services: resource discovery (using our mediation techniques) and optimistic replication (using our semantic reconciliation techniques).

7.3. Programme blanc ANR FLFS (2006-2009)

Participants: Jean Bézivin, Frédéric Jouault.

FLFS means "Families of Languages for Families of Systems" The objective of the FLFS project (with Obasco, Phoenix, INRIA, ENST) is to study the continuum domain modeling/implementation with the complementary technologies of model engineering, domain specific languages and aspect oriented programming. In this project, the Atlas team develops concrete solutions based on the AMMA platform and more specifically on the ATL and TCS languages.

7.4. RNTL OpenEmbeDD (2006-2009)

Participants: Jean Bézivin, Frédéric Jouault, Patrick Valduriez.

The project involves Airbus, Anyware, CEA, CS, FT R&D, LAAS, Thales DAE, Thales TRT and Verimag. OpenEmbeDD is an Eclipse open-source platform based on the model engineering principles for the software engineering of embedded systems. In this project, Atlas provides a model transformation virtual machine, and components for global model management. Furthermore, we address the interoperability between the AMMA and GME platforms.

7.5. RNTL XWiki Concerto (2006-2008)

Participants: Reza Akbarinia, Esther Pacitti, Mounir Tlili, Patrick Valduriez.

The project involves XPertNet, ObjectWeb, INRIA, ENST, Mandriva, and EISTI. The goal of the project is to enable Xwiki, an open source second generation wiki product, to operate in a P2P environment and support mobile users. In this project, Atlas develops with the ECOO INRIA project-team the technologies for collaborative editing of wiki documents in P2P.

7.6. OpenDevFactory, Paris Competitivity Cluster "Usine Logicielle" (2006-2008)

Participants: Freddy Allilaire, Jean Bézivin.

In this project, Atlas works, in relation with different industrial partners (Thales, EADS, Dassault, Esterel, Softeam, etc.) to develop model driven solutions to several software production problems. Based on our generic open source AMMA platform, complex chains of data transformations are considered as an alternative to more conventional development solutions.

8. Other Grants and Activities

8.1. Regional Actions

We are involved in two projects:

8.1.1. MILES (2007–2010)

N. Mouaddib coordinates the MILES project, funded by Region Pays-de-la-Loire. MILES is the main Region-funded project on information and communication technologies. Within the MILES project, M. Gelgon is in charge of a sub-project dealing with distributed multimedia systems, involving the Atlas project-team and IRCCyN (IVC group). This sub-project addresses, on one side, multimedia data learning and classification in a distributed computing and storage context and, on the other side, secure, distributed storage with involving techniques specific to multimedia data.

8.1.2. Pôle de compétitivité (2006-2008)

The ANR Safimage project (described further down) is further supported by Pôle de Competitivité Images & Réseaux.

8.2. National Actions

We are involved in four projects:

8.2.1. ARA Massive Data Respire (2006-2008)

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

The project involves LIP6 (leader), Paris (IRISA), Regal (INRIA et LIP6) and INT. The objective is to propose a P2P infrastructure for resource and data sharing in large scale networks. In this project, we study the following problems: resource catalog, dynamic clustering of peers, replication, query processing and equitable mediation. To validate the infrastructure, we develop services in the context of the APPA prototype.

8.2.2. ANR Safimage (2007-2010)

Participants: Marc Gelgon, Pierrick Bruneau.

This project involves Alcatel, IRCCyN, and IST. The project deals with inspection of data in high-speed routers for security purposes. The task devoted to Atlas is classification of multimedia data (examining how to scale up learning and recognition tasks with state-of-the-art classifiers in future routers).

8.2.3. ARA Massive Data SemWeb (2004-2007)

Participants: José Martinez, Noureddine Mouaddib, Guillaume Raschia.

The project SemWeb (Querying the Semantic Web with XQuery) 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.

8.2.4. ARA Massive Data APMD (2004-2007)

Participants: José Martinez, Noureddine Mouaddib, Guillaume Raschia.

The project APMD (Personalised Access to Masses of Data)(2004-2007) 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 through 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.

8.3. International actions

We are involved in the following international actions:

- 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), 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 contributes to the MDA work.

Furthermore, we have regular scientific relationships with research laboratories in

- North America: Univ. of Waterloo (Tamer Özsu), NYU (Dennis Shasha), New Jersey Institute of Technology (Vincent Oria), Wayne State University (Farshad Foutouhi and Wiliam Grosky), Kettering University (Peter Stanchev), Riad Hammoud (Delphi);
- Europe: Univ. of Madrid (Ricardo Jimenez-Periz), Univ. of Twente (Mehmet Aksit), Univ. of Roskilde (Henrik Larsen), Nokia (Andreas Myka), ;
- Others: Univ. Federal of Rio de Janeiro (Marta Mattoso), PUCPR, Curitiba, Brazil (Vidal Martins), Tokyo Metropolitan University (Hiroshi Ishikawa)

9. Dissemination

9.1. Animation of the scientific community

The members of the Atlas project-team have always been strongly involved in organizing the French database research community, in the context of the I3 GDR and the conference Bases de Données Avancées (BDA).

J. Bézivin is a member and co-founder of the steering committee of the ECOOP (AITO) and UML/Models conferences. In 2007, he was a co-chair of a track on model transformation of the ACM Symposium of Applied Computing.

In 2008, the Atlas project-team organizes the EDBT conference in Nantes. P. Valduriez is general chair; N. Mouaddib is executive chair; E. Pacitti is communication chair and G. Raschia is organization chair. G. Raschia is also co-chair of the Ph.D. workshop co-located with EDBT 2008. In 2009, P. Valduriez will be general chair of the VLDB conference in Lyon, with S. Abiteboul as PC chair.

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, Hermés : 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 :

- ACM-SIGMOD Int. Conf. 2007: P. Valduriez; 2008: P. Valduriez.
- IEEE Int. Conf. on Data Engineering (ICDE) 2007: P. Valduriez (Ind. PC chair), E. Pacitti.
- ACM Int. Conf. on Information and Knowledge Management (CIKM) 2007: E. Pacitti.
- European Dependable Computing Conference (EDCC) 2007: E. Pacitti.
- Int. Conf. on Very Large Databases (VLDB) 2008: P. Valduriez.
- IEEE Int. Conf. on Distributed Computing Systems (ICDCS) 2007, Data Management track: P. Valduriez (PC chair), E. Pacitti.
- Int. Conf. on High Performance Computing for Computational Science (VecPar) 2008: P. Valduriez.
- Workshop on P2P data management, co-located with EDBT 2008: E. Pacitti (PC co-chair).
- Int. Workshop on High-Performance Data Management in Grid Environments (HPDGrid) 2008, co-located with VecPar 2008: E. Pacitti (General chair), P. Valduriez (exec. chair).
- Int. Conf. on Advanced Information Systems Engineering (CAiSE) 2008: P. Valduriez.
- Journées Bases de Données Avancées (BDA), 2007: M. Gelgon, P. Valduriez.
- Int. Conf. on Extending DataBase Technologies (EDBT) 2008: E. Pacitti, G. Raschia.
- Enterprise Distributed Object Computing (EDOC), 2007: J. Bézivin.

9.3. Invited Talks

In April, P. Valduriez visited the university of Washington, Seattle (Prof. D. Suciu), for a week and gave a talk on data currency in DHTs. In April, E. Pacitti and P. Valduriez visited the university of Madrid (Prof. R. Jimenez-Péris), for a week and gave talks on data currency and replication in P2P.

In June, G. Raschia visited UCSB (Prof. A. el Abaddi, head of the CS department and leader of the Distributed System Lab.) and gave talks on data reduction techniques and data anonymization.

In July, E. Pacitti and P. Valduriez gave invited talks on data currency in DHTs and top-k query processing, respectively, at UFRJ, Rio de Janeiro and PUCPR, Curitiba, Brazil.

E. Pacitti gave an invited talk on replication in large-scale distributed systems at the seminar "30-year Perspective on Replication" organized by G. Alonso (ETH), B. Charron-Bost (Ecole Polytechnique), F. Pedone (Univ. of Lugano) and A. Schiper (EPFL) in Monte Verita, Switzerland in November.

9.4. Teaching

All the members of the Atlas project-team teach database management, multimedia, and software engineering at the Bs, Ms and Ph.D. degree level at the University of Nantes. José Martinez heads the computer science department at Polytech'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] R. AKBARINIA, V. MARTINS, E. PACITTI, P. VALDURIEZ. *Design and Implementation of Atlas P2P Architecture*, in "Global Data Management", R. BALDONI, G. CORTESE, F. DAVIDE (editors), IOS Press, 2006.

[2] R. AKBARINIA, E. PACITTI, P. VALDURIEZ. *Best Position Algorithms for Top-k Queries*, in "Int. Conf. on Very Large Data Bases (VDLB), Vienna, Austria", 2007, p. 495-506.

[3] R. AKBARINIA, E. PACITTI, P. VALDURIEZ. *Data currency in replicated DHTs*, in "ACM SIGMOD Int. Conf. on Management of Data (SIGMOD, Beijing, China", 2007, p. 211-222.

[4] F. ALLILAIRE, J. BÉZIVIN, F. JOUAULT, I. KURTEV. *ATL: a Model Transformation Tool*, in "Science of Computer Programming", to appear, 2007.

[5] I. KURTEV, K. VAN DEN BERG, F. JOUAULT. *Rule-based Modularization in Model Transformation Languages illustrated with ATL*, in "Science of Computer Programming", vol. 68, n^{o} 3, 2007, p. 111-127.

[6] A. NIKSERESHT, M. GELGON. *Gossip-based computation of a Gaussian mixture model for distributed multimedia indexing*, in "IEEE Transactions on Multimedia", to appear, 2007.

[7] E. PACITTI, P. VALDURIEZ, M. MATTOSO. *Grid Data Management: Open Problems and New Issues*, in "Journal of Grid Computing", vol. 5, n^{o} 3, 2007, p. 273-281.

[8] A. PIGEAU, M. GELGON. *Building and Tracking Hierarchical Partitions of Image Collections on Mobile Devices*, in "ACM Multimedia Conf., Singapore", 2005, p. 141-150.

[9] J. QUIANÉ-RUIZ, P. LAMARRE, P. VALDURIEZ. *SQLB: A Query Allocation Framework for Autonomous Consumers and Providers*, in "Int. Conf. on Very Large Data Bases (VDLB), Vienna, Austria", 2007, p. 974-985.

[10] R. SAINT-PAUL, G. RASCHIA, N. MOUADDIB. *General Purpose Database Summarization*, in "Int. Conf. on Very Large Databases (VLDB), Trondheim, Norway", 2005, p. 733-744.

[11] W. A. VOGLOZIN, G. RASCHIA, L. UGHETTO, N. MOUADDIB. *From User Requirements to Evaluation Strategies of Flexible Queries in Databases*, in "Handbook of Research on Fuzzy Information Processing in Databases", J. GALINDO (editor), to appear, Hershey, PA, USA, 2007.

[12] T. ÖZSU, P. VALDURIEZ. *Principles of Distributed Database Systems, 2nd edition*, Prentice Hall, 1999.

Year Publications

Doctoral dissertations and Habilitation theses

[13] R. AKBARINIA. *Techniques d'accès aux données dans les systèmes pair-à-pair*, Ph. D. Thesis, Université de Nantes, 2007.

[14] M. DIDONET DEL FABRO. *Gestion de métadonnées utilisant tissage et transformation de modèles*, Ph. D. Thesis, Université de Nantes, 2007.

[15] D. FAYE. *Médiation de données sémantique dans SenPeer, un système pair-à-pair de gestion de données*, Ph. D. Thesis, Université de Nantes, 2007.

[16] M. GELGON. *Structuration statistique de données multimédia pour la recherche d'information*, Habilitation Thesis, Université de Nantes, 2007.

[17] S. LEMP. *Médiation flexible dans un système pair-à-pair*, Ph. D. Thesis, Université de Nantes, 2007.

[18] V. MARTINS. *Réplication de données dans les systèmes pair-à-pair*, Ph. D. Thesis, Université de Nantes, 2007.

[19] W. A. VOGLOZIN. *Le résumé linguistique de données structurées comme support pour l'interrogation*, Ph. D. Thesis, Université de Nantes, 2007.

Articles in refereed journals and book chapters

[20] F. ALLILAIRE, J. BÉZIVIN, F. JOUAULT, I. KURTEV. *ATL: a Model Transformation Tool*, in "Science of Computer Programming", to appear, 2007.

[21] C. FURTADO, A. LIMA, E. PACITTI, P. VALDURIEZ, M. MATTOSO. *Adaptive Hybrid Partitioning for OLAP Query Processing in a Database Cluster*, in "Int. Journal of High Performance Computing and Networking", Special Issue on Best Papers from SBAC2005, in press, vol. 5, n^{o} 4, 2007.

[22] S. GANÇARSKI, H. NAACKE, E. PACITTI, P. VALDURIEZ. *The Leganet System: Freshness-aware Transaction Routing in a Database Cluster*, in "Information Systems", vol. 32, n^{o} 2, 2007, p. 320-343.

[23] F. JOUAULT, I. KURTEV. *On the Interoperability of Model-to-Model Transformation Languages*, in "Science of Computer Programming", to appear, Special Issue on Model Transformation, 2007.

[24] N. KOTOWSKI, A. A. B. LIMA, E. PACITTI, P. VALDURIEZ, M. MATTOSO. *OLAP Query Processing in Grids*, in "Concurrency and Computation: Practice and Experience", to appear, 2007.

[25] I. KURTEV, K. VAN DEN BERG, F. JOUAULT. *Rule-based Modularization in Model Transformation Languages illustrated with ATL*, in "Science of Computer Programming", vol. 68, n^{o} 3, 2007, p. 111-127.

[26] P. LAMARRE, S. LEMP, S. CAZALENS, P. VALDURIEZ. *A Flexible Mediation Process for Large Distributed Information Systems*, in "Int. Journal of Cooperative Information Systems", vol. 16, n^{o} 2, 2007, p. 299-332.

[27] Y. LIN, J. GRAY, F. JOUAULT. *DSMDiff: A Differentiation Tool for Domain-Specific Models*, in "European Journal of Information Systems", to appear, 2007.

[28] J. MARTINEZ. *L'indexation multimédia : description et recherche automatiques*, in "Formes et stratégies de l'interrogation de l'information visuelle (Traité IC2)", P. GROS (editor), Traitement du signal et de l'image, Hermes, Lavoisier, 2007.

[29] A. NIKSERESHT, M. GELGON. *Gossip-based computation of a Gaussian mixture model for distributed multimedia indexing*, in "IEEE Transactions on Multimedia", to appear, 2008.

[30] E. PACITTI, P. VALDURIEZ, M. MATTOSO. *Grid Data Management: Open Problems and New Issues*, in "Journal of Grid Computing", vol. 5, n^{o} 3, 2007, p. 273-281.

[31] J. E. ROUGUI, M. GELGON, D. ABOUTAJDINE, N. MOUADDIB, M. RZIZA. *Organizing Gaussian mixture models into a tree for scaling up speaker retrieval*, in "Pattern Recognition Letters", vol. 28, n^{o} 11, 2007, p. 1314-1319.

[32] W. A. VOGLOZIN, G. RASCHIA, L. UGHETTO, N. MOUADDIB. *From User Requirements to Evaluation Strategies of Flexible Queries in Databases*, in "Handbook of Research on Fuzzy Information Processing in Databases", J. GALINDO (editor), To appear, Hershey, PA, USA, 2008.

Publications in Conferences and Workshops

[33] R. AKBARINIA, E. PACITTI, P. VALDURIEZ. *Best Position Algorithms for Top-k Queries*, in "Int. Conf. on Very Large Data Bases (VDLB), Vienna, Austria", 2007, p. 495-506.

[34] R. AKBARINIA, E. PACITTI, P. VALDURIEZ. *Data currency in replicated DHTs*, in "ACM SIGMOD Int. Conf. on Management of Data (SIGMOD), Beijing, China", 2007, p. 211-222.

[35] R. AKBARINIA, E. PACITTI, P. VALDURIEZ. *Processing Top-k Queries in Distributed Hash Tables*, in "European Conf. on Parallel Computing (Euro-Par), Rennes", 2007, p. 489-502.

[36] R. AKBARINIA, E. PACITTI, P. VALDURIEZ. *Supporting Data Currency in Replicated DHTs using a Key-based Timestamping Service*, in "Journées Bases de Données Avancées (BDA), Marseille, France", 2007.

[37] M. BARBERO, F. JOUAULT, J. GRAY, J. BÉZIVIN. *A Practical Approach to Model Extension*, in "European Conf. on Model Driven Architecture- Foundations and Applications (ECMDA-FA), Haifa, Israel", LNCS, n^{o} 4530, Springer, 2007, p. 32-42.

[38] M. BECHCHI, G. RASCHIA, N. MOUADDIB. *Merging Distributed Database Summaries*, in "ACM Int. Conf. on Information and Knowledge Management (CIKM), Lisbon, Portugal", 2007, p. 419-428.

[39] M. DIDONET DEL FABRO, P. VALDURIEZ. *Semi-automatic Model Integration using Matching Transformations and Weaving Models*, in "ACM Symposium on Applied Computing (SAC), Seoul, Korea", 2007, p. 963-970.

[40] M. EL DICK, V. MARTINS, E. PACITTI. *A Topology-Aware Approach for Distributed Data Reconciliation in P2P Networks*, in "European Conf. on Parallel Computing (Euro-Par), Rennes", 2007, p. 318-327.

[41] M. EL DICK, E. PACITTI, P. VALDURIEZ. *Location-aware index caching and searching for P2P systems*, in "VLDB Int. Workshop on Databases, Information Systems, and Peer-to-Peer Computing (DBISP2P) - revised selected papers, Vienna, Austria", LNCS, to appear, Springer, 2007.

[42] D. FAYE, G. NACHOUKI, P. VALDURIEZ. *Semantic Query Routing in SenPeer, a P2P Data Management System*, in "Int. Conf. on Network-Based Information Systems (NBiS), Regensburg, Germany", 2007, p. 365-374.

[43] M. GELGON, A. NIKSERESHT. *Decentralized learning of a Gaussian Mixture with variational Bayes-based aggregation*, in "Euromicro Int. Conf. on Parallel and Distributed Processing, Toulouse, France", to appear, 2008.

[44] R. HAYEK, G. RASCHIA, P. VALDURIEZ, N. MOUADDIB. *Design of PeerSum: A Summary Service for P2P Applications*, in "Int. Conf. on Advances in Grid and Pervasive Computing (GPC'2007), Paris, France", 2007, p. 13-26.

[45] R. HAYEK, G. RASCHIA, P. VALDURIEZ, N. MOUADDIB. *Summary Management in P2P Systems*, in "Journées Bases de Données Avancées (BDA), Marseille, France", 2007.

[46] N. KOTOWSKI, A. A. B. LIMA, E. PACITTI, P. VALDURIEZ, M. MATTOSO. *OLAP Query Processing in Grids*, in "VLDB Int. Workshop on Data Management in Grids, Vienna, Austria", 2007.

[47] P. LAMARRE, J. QUIANÉ-RUIZ, P. VALDURIEZ. *Libra: une Méthode de Médiation Auto-Adaptative en Fonction des Attentes des Participants*, in "Journées Francophones sur les Systèmes Multi-Agents (JFSMA)", ISBN 2.85428.796.7, CEPADUES, 2007, p. 65-74.

[48] J. MANJARREZ-SANCHEZ, J. MARTINEZ, P. VALDURIEZ. *A Data Allocation Method for Efficient Content-Based Retrieval in Parallel Multimedia Databases*, in "Frontiers of High Performance Computing and Networking ISPA Workshops, Niagara Falls, Canada", LNCS, vol. 4743, Springer, 2007, p. 285-294.

[49] J. R. MANJARREZ-SANCHEZ, J. MARTINEZ, P. VALDURIEZ. *On the Usage of Clustering for Content Based Image Retrieval*, in "Int. Symp. on Computer Science in Russia (CSR), Ekaterinburg, Russia", LNCS, vol. 4649, Springer, 2007, p. 281-289.

[50] A. PIGEAU, A. NIKSERESHT, M. GELGON. *Fast tracking of hierarchical partitions with approximate KL-divergence for geo-temporal organization of personal images*, in "ACM Symp. on Applied Computing (SAC'2007), Seoul, Korea", 2007, p. 1088-1089.

[51] J. QUIANÉ-RUIZ, P. LAMARRE, S. CAZALENS, P. VALDURIEZ. *A Satisfaction Balanced Query Allocation Process for Distributed Information Systems*, in "Journées Bases de Données Avancées (BDA), Marseille, France", 2007.

[52] J. QUIANÉ-RUIZ, P. LAMARRE, S. CAZALENS, P. VALDURIEZ. *Satisfaction Balanced Mediation*, in "Int. Conf. on Information and Knowledge Management (CIKM), Lisbon, Portugal", 2007, p. 947-950.

[53] J. QUIANÉ-RUIZ, P. LAMARRE, P. VALDURIEZ. *KnBest - A Balanced Request Allocation Method for Distributed Information Systems*, in "Int. Conf. on Database Systems for Advanced Applications (DASFAA), Bangkok, Thailand", 2007, p. 237-248.

[54] J. QUIANÉ-RUIZ, P. LAMARRE, P. VALDURIEZ. *SQLB: A Query Allocation Framework for Autonomous Consumers and Providers*, in "Int. Conf. on Very Large Data Bases (VDLB), Vienna, Austria", 2007, p. 974-985.

[55] J. QUIANÉ-RUIZ, P. LAMARRE, P. VALDURIEZ. *Un Modèle pour Caractériser des Participants Autonomes dans un Processus de Médiation*, in "Modèles Formels de l'Interaction (MFI)", ISSN 1762-455X, LAMSADE, 2007, p. 389-396.

[56] R. SAINT-PAUL, G. RASCHIA, N. MOUADDIB. *Database Summarization: The SaintEtiQ System*, in "IEEE Int. Conf. on Data Engineering (ICDE), Istanbul, Turkey", 2007, p. 1475-1476.

[57] A. VENTRESQUE, S. CAZALENS, P. LAMARRE, P. VALDURIEZ. *Query Expansion And Interpretation To Go Beyond Semantic Interoperability*, in "Int. Conf. on Ontologies, DataBases, and Applications of Semantics (ODBASE), Algarve, Portugal", LNCS, Springer, 2007, p. 870-877.

Miscellaneous

[58] P. BRUNEAU, M. GELGON, F. PICAROUGNE. *Incremental, semi-supervised classification for surveillance of multimedia content in future high-speed routers*, 2007, NATO Advanced Study Institute on Mining Massive Data Sets for Security, poster.

[59] E. E. SENTURK. *Sélection à la volée du niveau de précision dans une hiérarchie de résumés de données*, 2007, rapport de Master, Université de Nantes.

[60] E. E. SENTURK. *Sélection à la volée du niveau de précision dans une hiérarchie de résumés de données*, Technical report, Université de Nantes, 2007.