INRIA

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** . . . . . **1**
**2. Overall Objectives** . . . . . **1**
2.1. Introduction 1
2.2. Highlights of the Year 2
**3. Scientific Foundations** . . . . . **2**
3.1. Data Management 2
3.2. Data Reduction Techniques 3
3.3. Distributed Data Management 4
3.4. Data Stream Management 5
3.5. Semantic Interoperability 6
**4. Application Domains** . . . . . **6**
**5. Software** . . . . . **7**
5.1. APPA (Atlas Peer-to-Peer Architecture) 7
5.2. KTS (Key-based Timestamp Service) 7
5.3. P2P-LTR (P2P Logging and Timestamping for Reconciliation) 8
5.4. SbQA (Satisfaction-based Query Allocation Framework) 8
5.5. PeerUnit (Peer-to-Peer Tester) 8
5.6. DBSum 8
**6. New Results** . . . . . **9**
6.1. Data Reduction and Classification 9
6.1.1. Database and Data Stream Summaries 9
6.1.2. Querying Database Summaries 9
6.1.3. Distributed Learning of Probabilistic Class Models 10
6.2. Data Access with Autonomous Participants 10
6.2.1. Satisfaction-based query allocation 11
6.2.2. Data Access in the Context of Heterogeneous Semantics 11
6.3. P2P Data Management 12
6.3.1. Management of Replicated Data 12
6.3.2. Data Clustering and Query Routing 13
6.3.3. Testing P2P Systems 13
6.4. Data Access using DHTs 14
6.4.1. Data streaming 14
6.4.2. Data Privacy 15
**7. Contracts and Grants with Industry** . . . . . **15**
7.1. STREP Grid4All (2006-2009) 15
7.2. RNTL XWiki Concerto (2006-2008) 15
7.3. ANR Safimage (2007-2010) 16
**8. Other Grants and Activities** . . . . . **16**
8.1. Regional Actions 16
8.1.1. MILES (2007–2010) 16
8.1.2. Pôle de compétitivité (2006-2008) 16
8.2. National Actions 16
8.2.1. ARA Massive Data Respire (2006-2008) 16
8.2.2. ARA Massive Data SemWeb (2004-2008) 16
8.2.3. ARA Massive Data APMD (2004-2008) 16
8.3. International actions 17
**9. Dissemination** . . . . . **17**
9.1. Animation of the scientific community 17
9.2. Editorial Program committees 17

9.3. Invited Talks 18
9.4. Teaching 18
**10. Bibliography** ..... **18**

# 1. Team

**Research Scientist**

Patrick Valduriez [ Team Leader, Research Director (DR) Inria, HdR ]

**Faculty Member**

Marc Gelgon [ Professor, HdR ]
Philippe Lamarre [ Associate Professor ]
Esther Pacitti [ Associate Professor, HdR ]
Guillaume Raschia [ Associate Professor ]
Gerson Sunyé [ Associate Professor ]

**Technical Staff**

Jorge Quiané Ruiz [ Engineer, Grid4All since oct. ]
Mounir Tlili [ Engineer, Xwiki Concerto until oct. ]

**PhD Student**

Eduardo Almadeia [ ALBAN fellowship until Sept., ATER U. Nantes since oct. ]
Mounir Bechchi [ ATER U. Nantes ]
Pierrick Bruneau [ ANR funding ]
William Kokou Dedzoe [ CNRS fellowship since Oct. ]
Fadi Draidi [ MAE fellowship since Oct. ]
Ali El Attar [ Pays-de-la-Loire fellowship ]
Rabab Hayek [ MENRT fellowship ]
Manal El Dick [ MENRT fellowship ]
Mohamed Jawad [ MENRT fellowship ]
Jorge Manjarrez Sanchez [ Conacyt fellowship ]
Afshin Nikeresht [ Iranian gov. fellowship, until Oct. ]
Wenceslao Palma [ INRIA fellowship ]
Quang-Khai Pham [ CNRS-U. of New South Wales fellowship ]
Jorge Quiané Ruiz [ Conacyt fellowship until Sept. ]
Jamal Rougui [ Egide funding, until July ]
Toufik Sarni [ CNRS fellowship since Oct. ]
Mounir Tlili [ MILES fellowship since Nov. ]
Anthony Ventresque [ ATER U. Nantes ]
Amenel Voglozin [ ATER U. Nantes until Sept. ]

**Administrative Assistant**

Élodie Lizé [ Secretary, CDD ]

# 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. Second, data management systems need to scale up to support large-scale distributed 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 now focus on data management in two large-scale distributed contexts: the web and P2P systems. In the context of the web, we consider information systems with autonomous participants (with heterogeneous data and different interests) and deal witht the problems of data integration, data classification and data access. In the context of P2P systems, we capitalize on our experience in developping the APPA system, with various data management services (replication, caching, queries, clustering, testing, etc.). A hard problem that we are starting to investigate is data privacy in P2P systems.

## 2.2. Highlights of the Year

The first highlight of the year is that Reza Akbarinia (postdoc at U. of Waterloo) obtained a CR2 position in Atlas; he will join us in january 2009. The second highlight is the spinoff of our activity on model management in the INRIA team AtlanMod created by Jean Bézivin (a previous member of the Atlas project-team) in july. Started within Atlas in 2003, the activity on model management has become highly visible in terms of research results, industrial contracts and technology transfer. And it has grown to the point that it will be best pursued within a separate, more focused team (see the AtlanMod 2008 activity report for details on its objectives). However, Atlas will continue to collaborate with AtlanMod on data engineering problems, in particular, in the context of data integration.

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

- a new ANR project (VERSO program) DataRing on P2P data sharing for online communities (headed by P. Valduriez, with the Gemo project-team, LIG, LIRMM and Telecom Paristech) for 3 years, starting in 2009;
- a very good publication record, including 3 papers published in top database journals (VLDBJ, IEEE TKDE, DAPD);
- the paper [34], in the context of the CNPq-INRIA Gridata project, obtained the best student paper award at VecPar08;
- the paper [45] which proposes a new technique of semantic query interpretation was nominated for best paper (top 4) at ESWC08;
- 1 HDR and 4 Ph.D. theses defended;
- the release of two new APPA services: P2P-LTR (logging and timestamping for reconciliation) which was demonstrated at several conferences (VLDB08, BDA08, NOTERE08); and Satisfaction-based Query Allocation Framework (SbQA) which will be demonstrated at ICDE09.

# 3. Scientific Foundations

## 3.1. Data Management

**Keywords:** *Data independence*, *consistency*, *metadata modelling*, *queries*, *schema management*, *transactions*.

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 independence*, 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, multimedia data and data streams, in highly distributed environments. Thus, we also capitalize on scientific foundations in data reduction techniques, distributed data management, data stream management and semantic interoperability to address these problems.

## 3.2. Data Reduction Techniques

**Keywords:** *Data aggregation*, *approximate anwering*, *data classification*, *database summaries*, *multimedia data*, *probabilistic models*.

With the explosion of the quantities of data to be analyzed, it is desirable to sacrifice the accuracy of the answers for response time. Particularly in the early, more exploratory, stages of data analysis, interactive response times are critical, while tolerance for approximation errors is quite high. In this context, data reduction is important to control the desired trade-off between answer accuracy and response time.

Data reduction is closely associated with aggregation. While histograms form the baseline approach and have been extensively used for query optimizers, a wealth of techniques have been proposed. In particular, cluster-based reduction of data, where each data item is identified by means of its cluster representative, leads to classical tree indexes, where data is partitioned recursively into buckets. The clusters may be data-driven, or independent from the data. With minimal augmentation, it becomes possible to answer queries approximately based upon an examination of only the top levels of an index tree. If these top levels are cached in memory, as is typically the case, then one can view these top levels of the tree as a reduced form of data suitable for approximate query answering.

To deal with large amounts of data, or high-dimensional data, much work has also been devoted to reducing the dimension of representations, by identifying lower dimension manifolds on which data essentially lies. Single-value decomposition or discrete wavelet transformations are two examples of such transform-based techniques. Among data reduction techniques, one may further distinguish parametric techniques (e.g. linear regression), that assume a model for the data, from non-parametric techniques. While the former offer generally more compression, automatically selecting the form of the model remains a difficult issue.

An important use of data reduction is for retrieval within collections of multimedia material, such as image, audio or video. For the purpose of comparing queries to target documents or for building an index, these documents are represented by features, i.e. multivariate attributes. These features may be used directly (e.g. nearest neighbourhood search among feature vectors, for image matching) or, often, through probabilistic models of

their distribution, thereby capturing the variability of a given class. The design of these features requires a specific expertise for each media, to ensure a good trade-off between concision, ability to discriminate and invariance to certain imaging or acoustic conditions. This is typically handled by media-specific research communities.

Nearest neighbour queries are appropriate for multimedia information retrieval. Efficient multimedia feature vectors often span high dimensional spaces, where indexing structures classically used in database management systems (tree-based and hashing-based) are not effective, due to the dimensionality curse. Parallel databases may contribute to maintaining reasonnable query processing time, but require the definition of data distribution strategies. Such strategies are one of the focuses of our work.

Among models, parametric probabilistic models build a very rich, well-founded and well-documented toolbox for representing the data distributions in a concise way, in association to statistical estimation techniques for determining the form of the model and values of its parameters. Together, they provide a strong share of existing solutions to multimedia data analysis problems (learning and recognition). Relating this to database summaries, seeking simple forms to *describe* the data (structure for efficient retrieval) and forms that *explain* the data (structure for understanding, where parametric forms introduce the necessary inductive biais) are often very close goals, hence a growing number of techniques common to the database and machine learning communities. Among probabilistic models, generative mixture models consider the data to be a combination of several populations, whether this correspond to true variety of natures or whether is a only a modelling tool. Mixtures have wide modelling ability, like non-parametric methods, but retain the parsimony of parametric approaches. Hence, they have been much studied, extended and applied, in the contexts of both supervised and unsupervised learning. In the case of probabilistic models, Bayesian estimation supplies a principled solution to the abovementionned model selection. This long remained either computation-intensive or very approximative, but nowadays, besides increasing computing power available, a corpus of efficient approximate inference mechanisms has been built, for a growing variety of graphical model structures. There remain questions which are receiving growing attention : how can such models be efficiently learned from dynamic distributed data sources ? How can a large set of probabilistic models be indexed ?

Among the broad range of reduction techniques, the database summarization paradigm has become an ubiquitous requirement for a variety of application environments, including corporate data warehouses, network-traffic monitoring and large socio-economic or demographic surveys. Besides, downsizing massive data sets allows to address some critical issues such as individual data obfuscation, optimization of the usage of system resources like storage space and network bandwidth, as well as effective approximate answers to queries. Depending on the application environment and the preferred goal of the approach, we distinguish three families of approaches concerned with database summarization. The first one focuses on aggregate computation and it is supported by statistical databases, OLAP cubes and multdimensional databases.The second class of approaches extends the previous one in that it tries to produce more compact representations of aggregates. The main challenge for such methods is to keep expressiveness of the provided access methods (aggregate queries) to the items without any need to uncompress the structure. Quotient cubes and linguistic summaries are two major contributions in that direction. The third family of approaches deals with intentional characterization of groups of individuals based on usual mining algorithms. Those categories are obviously not sharp and there are many orthogonal criteria that encompass such a classification. For instance, some of them share the same theoretical background (Zadeh's fuzzy set theory) and they use fuzzy partitions and linguistic variables to support a robust summarization process.

This database research field raises new challenges, in particular, to push more semantics into summaries while still remaining efficient in the context of database systems. Update of such metadata is also of major concern. Furthermore, traditional problems of data management such as query evaluation or data integration have to be revisited from the point of view of database summaries.

## 3.3. Distributed Data Management

**Keywords:** *Distributed systems*, *P2P systems*, *caching*, *data integration*, *replication*, *top-k query processing*.

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) [10]. 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.

## 3.4. Data Stream Management

**Keywords:** *Data stream*, *approximate answering*, *continuous query*, *sliding window*.

Recent years have witnessed major research interest in data stream management systems. A data stream is a continuous and unbounded sequence of data items. There are many applications that generate streams of data including financial applications, network monitoring, telecommunication data management, sensor networks, etc. Processing a query over a data stream involves running the query continuously over the data stream and generating a new answer each time a new data item arrives. Due to the unbounded nature of data streams, it is not possible to store the data entirely in a bounded memory. This makes difficult the processing of queries that need to compare each new arriving data with past ones. A common solution to the problem of processing join queries over data streams is to execute the query over a sliding window that maintains a restricted number of recent data items. This allows queries to be executed in a finite memory and in an incremental manner by generating new answers when a new data item arrives. Due to the continuous, often very fast, arrival of

new data, it is impossible to produce exact answers to queries. Therefore, approximate answers are typically provided.

In real data settings, a data stream management system may process hundreds of user queries. Therefore, for most realistic distributed streaming applications the naive solution of collecting all the data at a single site is simply not viable. Therefore, we are interested in techniques for processing continuous queries over collections of distributed data streams. An example of such queries is join queries which are very important for many applications. A streaming join computation can be useful in understanding important trends and making decisions about measurements or utilization patterns.

## 3.5. Semantic Interoperability

**Keywords:** *Semantic web*, *ontology alignement*, *semantic representations*.

Semantic interoperability ensures that the meaning of the information that is exchanged is automatically interpreted by the receiver of a message. In centralized systems, this property improves the relevance of query answers. In distributed heterogeneous systems, it is compulsory to enable autonomous heterogeneous sources understand each other to obtain relevant results.

To provide semantic interoperability within a system, much research has been conducted on semantic representations. The main idea is to use meta-information which eases the meaning understanding. This approach needs the definition of ontologies which describe the concepts and relations between them, for a given domain. During the last fifteen years, much effort has focused on formal methods to describe ontologies, ressource description languages, reasoning engines...All these methods represent the fundations of the semantic web. However, many works rely on the assumption that a single ontology is shared by all the participants of the system.

However, in distributed systems with autonomous participants, such as P2P systems, this assumption is not realistic anymore. On the contrary, one has to consider that the participants create their ontologies independently of each other. Thus, most often the ontologies differ. To tackle this problem, research on ontology matching proposes several techniques to define corresondances between entities of two ontologies. So, in some way, ontology matching highlights the shared parts of two ontologies. Thus it provides the basis for interoperability between heterogeneous participants and by “transitivity” in the whole system.

Although ontology matching and other semantic web techniques provide a basis for interoperability, the challenge is still to define a whole semantic infrastructure in which participants’ search for information is both relevant and efficient. Considering semantics can be useful at different stages. First, semantic representation of queries and information may improve the relevance of the results. It can be used in place or in addition to usual request representation. Second, semantics can be used to represent participants, or groups of them, leading participants to better know each other. Such information can be useful for routing the requests to other participants in order to obtain the relevant answers within a short time and with a low trafic load. Third, this information can also be used to organize the network so as to improve efficiency. All these research directions have received partial answers but more work is needed on the interaction between all these elements and their impact on the efficiency of the global system.

# 4. Application Domains

## 4.1. Overview

**Keywords:** *Distributed collaborative applications*, *decision-support applications*, *multimedia personal databases*, *professional communities*, *social networks*.

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. Other examples of collaborative applications we are interested in are social networks. In Atlas, we address the issues of data sharing 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. APPA (Atlas Peer-to-Peer Architecture)

**Participants:** Eduardo Almadeia, William Kokou Desdoe, Philippe Lamarre, Esther Pacitti, Gerson Sunyé, Jorge Quiane Ruiz, 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, OpenChord and Pastry and tested on GRID5000 and PlanetLab. The current services of APPA are (see below): KTS, SbQA, P2P-LTR and PeerUnit. The APPA services are used in several projects: Strep Grid4All, ANR RNTL Xwiki Concerto and ANR VERSO DataRing.

## 5.2. KTS (Key-based Timestamp Service)

**Participants:** William Kokou Desdoe, Esther Pacitti, Patrick Valduriez [contact].

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

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. KTS has been initially proposed to support data currency in DHTs, i.e. the ability to return a current replica in a DHT despite peers leaving the network or concurrent updates. Experimental validation has shown that KTS incurs very little overhead in terms of communication cost. KTS is the basis for the P2P-LTR service. It has been implemented in Java on top of OpenChord.

## 5.3. P2P-LTR (P2P Logging and Timestamping for Reconciliation)

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

URL: http://p2pltr.gforge.inria.fr/

P2P-LTR provides two major functions: logging of user actions in a DHT and continuous, distributed timestamping of these actions. This is useful to perform reconciliation of replicated data. P2P-LTR extends KTS with continuous timestamping and logging of actions. To perform reconciliation using P2P-LTR, we use a simple reconciliation algorithm based on operational transforms, called SB, from the ECOO team at LORIA and readily available as Open Source Software. P2P-LTR has been implemented in Java on top of OpenChord. It has been validated in the Strep Grid4All and RNTL Xwiki Concerto projects to perform reconciliation of replicated documents in a P2P wiki system.

## 5.4. SbQA (Satisfaction-based Query Allocation Framework)

**Participants:** Philippe Lamarre [contact], Jorge Quiane Ruiz, Patrick Valduriez.

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

SbQA is a Satisfaction-based Query Allocation framework for distributed environments where consumers and providers are autonomous and have special interests towards providers and queries, respectively. We experimentally demonstrated that it ensures good system performances while satisfying consumers and providers. Hence, SbQA can scale-up in these environments by preserving the total system capacity, i.e. the aggregate capacity of all providers. SbQA is used in the Strep Grid4All project as the basis to perform selection of services proposed by market-places as well as altruist contributors. SbQA is implemented in Java.

## 5.5. PeerUnit (Peer-to-Peer Tester)

**Participants:** Eduardo Almadeia [contact], Gerson Sunyé, Patrick Valduriez.

URL: http://peerunit.gforge.inria.fr/

Peerunit is a testing framework for P2P systems. It is useful to developers who want to implement unit tests for a Java P2P system. The framework is based on two original aspects: (i) the individual control of peers volatility and (ii) a distributed testing architecture to cope with large numbers of peers. A distributed component, the tester, executes on peers, and controls their execution and their volatility, making them leave and join the system at any time, according to the needs of a test. Furthermore, testers communicate with each other across a balanced tree (B-Tree) structure to avoid using a centralized testing coordination. Peerunit is implemented in Java and ha sbeen validated on two popular open-source P2P systems (FreePastry and OpenChord).

## 5.6. DBSum

**Participants:** Mounir Bechchi, 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. Data Reduction and Classification

Data reduction and classification is needed to cluster large data sets in concise ways. We use two different formalisms for clustering data: grid-based conceptual hierarchies, for database summarization; and parametric probabilistic models, for continuous multivariate spaces typically encountered with multimedia data. To deal with distributed data sources, we have addressed the problem of integration of (possibly hierarchical) structures. Our focus is on integration of data descriptions, without resorting to raw data. We have also addressed the problem of efficient querying of database summaries.

### *6.1.1. Database and Data Stream Summaries*

**Participants:** Guillaume Raschia, Mounir Bechchi, Quang-Khai Pham.

Our database summarization system DBSum provides multi-level summaries of tabular data stored in a centralized database. Summaries are computed online by means of a grid-based conceptual hierarchical clustering algorithm. Along this research direction, we addressed two problems related to data reduction: (i) joining two summaries from a database join operation of raw data, and (ii) defining summaries of temporal sequence data.

We then proposed [60] two new algorithms for summarizing heterogeneous, distributed data without a prior unification of the data sources in order to overcome the limitations of the centralized version: Subspace-Oriented Join Algorithm (SOJA) and Tree Alignement-based Join Algorithm (TAJA). The main idea of such algorithms consists in applying innovative join operators on two local models, computed over two disjoint sets of features, to provide a global summary over the full feature set without scanning the raw data. An experimental study have shown that our joining processes (SOJA and TAJA) yield high quality clustering of the entire distributed data and are very efficient in terms of computational time w.r.t. the centralized approach.

As a second major contribution to data reduction, we proposed a Time-Aware Content Summarization (TACS) technique for large historical data sources [37]. The TACS process takes as input user background knowledge in the form of taxonomies and a historical data source where items are ordered following increasing timestamps. It produces a reduced version of the data through a Generalize & Merge process where items are expressed at higher levels of abstraction while preserving their temporal ordering. We formalized the summarization features such as domain and numerosity reduction capabilities, time-awareness and described its incremental and reduced complexity aspects. The resulting output is a representation of the data in a more reduced and concise form, yet informative enough to support time-sensitive applications such as Sequential Pattern Mining (SPM). We applied SPM over TACS summaries, demonstrated the benefits of the summary for such application and validated our study with experimental results on a real world data set from Reuters news archives.

### *6.1.2. Querying Database Summaries*

**Participants:** Amenel Voglozin, Guillaume Raschia, Mounir Bechchi.

We investigated the effective usage of database summaries both for the many-answsers problem and as an underlying structure for the evaluation of fuzzy queries. In typical querying within large databases, the user formulates a first basic (broad) query to target and filter data and next, she starts browsing the answer looking for precise information. To quickly provide the user with concise, useful and structured answers as a starting point for online exploration, we proposed to perform an offline hierarchical grid-based clustering of the data set [61]. Every single answer item describes a subset of the queried data in a user-friendly form using linguistic labels, i.e. it represents a concept that exists within the data. Furthermore, the answers of a given *blind query* are nodes of a classification tree and every subtree rooted by an answer offers a *guided tour* of a data subset to the user. Our solution is able to deal with fuzzy predicates in the query such that the user's perception of the domain is observed.

We also addressed the problem of fuzzy querying. In [55], we discuss the whole process of fuzzy querying, from the query formulation to its evaluation. Mainly, it advocates the use of index structures in the evaluation of fuzzy queries and it presents the DBSum-based evaluation algorithm as an alternative plan for such queries. In [24], we addressed the problem of free vocabulary predicates in fuzzy queries, i.e. the way end-users could formulate a personalized query against database summaries.

### *6.1.3. Distributed Learning of Probabilistic Class Models*

**Participants:** Pierrick Bruneau, Ali El Attar, Marc Gelgon, Afshin Nikeresht.

Learning a probabilistic model that describes the distribution of numerical features in a multidimensional continous space, for supervised or unsupervised classification, is a fundamental and widely studied task. For instance, it takes a central place in multimedia information retrieval. When data sources are distributed and dynamic, existing solutions must be revisited. We indeed foresee a growing importance of this research in multimedia learning and recognition.

Our focus had been on mixture model aggregation that operates on mixture model parameters rather than data [22], with a view to modelling a single class density. A first open issue is the definition of a scheme for determining efficiently the suitable number of components in the aggregated mixture. We showed in [26] how this could be conducted by means of a variational-Bayes approach. We applied this technique for mixture model estimation in the context of gossip propagation of parameters, to suit the context of dynamic distributed sources [31].

A variant of this work aims at data clustering. As a side product to the abovementioned technique, we described in [27] how the task of multiscale clustering may benefit from the scheme. However, the use of Gaussian mixture for clustering is known to suffer from the lack of statistical robustness. We are currently devising a solution for distributed clustering, founded on the same general approach, but based on mixtures of $t-$distributions. The challenge lies in the lack of obvious closed-form solution for parameter-based mixture aggregation and, probably, the need to define an accurate yet low-cost approximative inference scheme. To this end, we are considering utilizing the decomposition of $t-$distribution as infinite gamma-weighted Gaussian mixture.

Finally, we are considering joint learning of multiple classes, where each class model is described by a mixture. More precisely, we are addressing the case where only some of the sources supply a class label along with their local model, or where contraints for class assignment are provided, e.g. two sources are known to provide models for the same or different class, but the identities of these classes are unknown. This comes down to semi-supervised learning, but operating on a set of models rather than pointwise data. While the main applicative scenario is that of class model learning from distributed sources, we are also using this work in a cooperation with the COD research team at LINA, for interactive classification with visualization [53], in the framework of the ANR Safimage project [64].

## 6.2. Data Access with Autonomous Participants

Taking into account the autonomy of participants holds an important part in the evolution of data management, systems and applications, especially considering open systems such as internet. Autonomy is some kind of

freedom left to participants which can be managed differently from one participant to another. It can take very different forms. For example, it may mean "enter or leave the system at will" as managed in P2P systems. Intuitively, the more participants are autonomous, the easier it is for a participant to integrate the system, but, the harder it will be to manage the system because participants'autonomy induces heterogeneity. In this context, we have focused on two dimensions of participants' autonomy: objectives' autonomy and semantic heterogeneity.

### *6.2.1. Satisfaction-based query allocation*

**Participants:** Philippe Lamarre, Jorge Quiane Ruiz, Patrick Valduriez.

Objectives' autonomy reflects participants' objectives in joining a system. Intuitively, in the field of open systems, it is the hope of achieving some objectives which motivates a peer to participate in a system. Obviously, different participants may have different objectives. Thus, to integrate as many participants as possible, a system should not assume that all of them are interested in the same normative objective. Instead, it should enable participants to act accordingly to their own private objectives. This approach has been studied in the field of query allocation.

In the context of dynamic distributed systems, with large numbers of heterogeneous, autonomous consumers and providers (the participants) query allocation is challenging because participants' interests may be contradictory. For example, a consumer would desire to receive results from a given provider but this provider would not desire to perform the query of such a consumer. In [23], we defined a model to characterize the participants' satisfaction in the long run. We proposed a query allocation framework that takes into account participants' interests so as to satisfy them in the long-run. An important aspect of our solution is that it dynamically balances between providers' interests and consumers' interests taking into account their respective rewards.

Our experimental results showed that our model enables a better evaluation of query allocation methods in these environments, and that our query allocation approach significantly outperforms baseline methods from both a satisfaction and performance point of view. The solution in [23] was proposed for mono-mediator systems. In [38], we proposed an economic version for multi-mediator systems. This solution uses "virtual money" as a means of regulation so as to allow several mediators to perform query allocation by preserving participants' satisfaction and system performance as in mono-mediator systems. A strong point of this solution is that its system regulation is independent of the number of mediators in the system, which allows to easily scale up. We compared our economic solution to some economic baseline methods and demonstrate its superiority. Our experimental results have demonstrated that our economic query allocation approach allows scaling up to several mediators without any loss in performance and participants' satisfaction.

In [40], we presented our query allocation prototype using a volunteer computing-based application (such as the BOINC platform) and demonstrated the efficiency of our solution to perform query allocation as well as its self-adaptability to participants' interests.

### *6.2.2. Data Access in the Context of Heterogeneous Semantics*

**Participants:** Philippe Lamarre, Anthony Ventresque, Patrick Valduriez.

The second dimension of autonomy we consider is related to semantic heterogeneity. The general objective is to provide each participant with a maximum of autonomy such that a participant's effort on the development of a semantical representation is not comprimised by joining the system. The main problem is then to enable correct communication between participants using many different semantic representations. Using mapping and alignments between ontologies, our goal is to go beyond the corresponding parts of the ontologies, thus enabling a participant to use her full ontology to express her requests and to answer. When this is possible, semantic autonomy becomes an asset for the system and an added value for those who develop semantic representations. This approach has been studied in the field of information retrieval.

We proposed a new three-step intuitive approach [44], [45], [46]. First, to make its query understandable, a requester must explain the concepts used (structured expansion). Second, when receiving a request, a provider, uses these explanations to express the requester's needs using its own concepts. Last but not least, to evaluate the relevance of documents, their representations are adapted with respect to the explanations given in the request. For example, if the request explanations consider the two concepts "cat" and "wildcat" to be synonyms, while according to the local ontology they are different, the document's representation needs to be modified in order to conform to requester representation.

These considerations led us to study different components of this problem, for example, we explored similarity measures between concepts within an ontology in [52]. We studied and evaluated our solution within a semantically homogeneous context. The conclusion is that our solution does not improve significantly the results quality obtained by the information retrievial process, but more important, does not degrade the results. Thus, our approach can be used in an homegeneous context. The most interesting results have been obtained in a semantically heterogeneous context where our approach improves significantly the information retrieval results [45]. For example, it allows to retrieve over 80% of relevant information which would have been retrieved if the context were homogeneous, where classical approches only retrieve 30%.

## 6.3. P2P Data Management

Data management in P2P systems offers new research opportunities since traditional distributed database techniques need to scale up while supporting data autonomy, heterogeneity, and dynamicity. In the context of the Atlas Peer-to-Peer Architecture (APPA) project, the main results this year are in the management of replicated data, data clustering and testing.

### *6.3.1. Management of Replicated Data*

**Participants:** William Kokou Desdoe, Manal El Dick, Esther Pacitti, Mounir Tlili, Patrick Valduriez.

Data replication is important to improve performance and provide fault-tolerance in P2P systems. We addressed data replication in two different application contexts: collaborative work (e.g. collaborative text editing) and web content caching. Large-scale distributed collaborative applications are now getting common as a result of rapid progress in distributed technologies (grid, P2P, and mobile computing). An example of such applications is a second generation Wiki that works over a P2P network and supports users on the elaboration and maintenance of shared documents in a collaborative and asynchronous manner. To provide scalability for those applications, optimistic multimaster replication is necessary so that the same document may be updated in parallel by different users. Then, the main problem is to manage P2P data reconciliation in the presence of updates on replicated documents. We addressed this problem in two separate solutions.

The first solution developped in [21] is a distributed algorithm for semantic reconciliation in P2P networks (P2P-reconciler). Other important contributions are: a basic cost model for computing communication costs in a DHT network; a strategy for computing the cost of each reconciliation step taking into account the cost model; and an algorithm that dynamically selects the best nodes for each reconciliation step. Furthermore, since P2P networks are built independently of the underlying topology, which may cause high latencies and performance penalties, we proposed a topology-aware variant of our P2P-reconciler algorithm and showed the important performance gains. This work was done in cooperation with V. Martins (now PUCPR, Curritiba) and R. Jimenez-Periz (Univ. Madrid).

The second solution proposed in [43], [42] is P2P-LTR (P2P Logging and Timestamping for Reconciliation) to perform distributed reconciliation over a DHT. It extends the Key-based Timestamping Service proposed in 2007 to support continuous, decentralized timestamping. While updating at collaborating peers, updates are timestamped and stored in a highly available P2P log. During reconciliation, these updates are retrieved in total order to enforce eventual consistency despite churn and failures. We implemented P2P-LTR using OpenChord and demonstrated it with various failure scenarios in the context of a wiki [28]. The demonstration shows decentralized, continuous timestamp management for concurrent updating, logging in the DHT for high availability of updates, and retrieval of timestamped updates from the P2P-Log in total order. This work was done in cooperation with Reza Akbarinia (postdoc, Univ. Waterloo) who will join us in 2009 as CR2 INRIA.

We also started to address P2P data replication for web content caching. The main motivation ofr using a P2P approach is that many websites with a large user base, e.g., non-profit organizations, do not have the financial means to install large web-servers or use specialized content distribution networks (CDN) such as Akamai. For those websites, we proposed Flower-CDN [62][29], a locality-aware P2P CDN whereby users interested in a website can provide support for distribution of its content. The main idea is that peers keep the web pages they retrieve and later serve them to other peers that are near them in the network. The architecture of Flower-CDN is hybrid between structured and unstructured networks. When a peer requests a web page from a website for the first time, a locality-aware DTH quickly finds a peer in its neighborhood that has the web page available. Additionally, all peers in a given region that maintain content of a particular website build an unstructured content overlay. Within a content overlay peers gossip information about their content, thus allowing the system to maintain accurate information despite failures and churn. This work was done in cooperation with Bettina Kemme (Mc Gill Univ.).

In the context of the CNPq-INRIA Gridata project (with UFRJ and PUC-Rio, Brazil), we also exploited replication to increase performance in large-scale data management systems. In [34], we validated the results of previous years with the ParGRES database cluster on high-performance query processing using a real-world OLAP database. In [51], we also demonstrated the effectiveness of replication together with data fragmentation to speed up the execution of BLAST (Basic Local Alignment Search Tool), a popular tool for (bio)sequences comparison and alignment in Genome sequencing and analysis.

### *6.3.2. Data Clustering and Query Routing*

**Participants:** Rabab Hayek, Guillaume Raschia, Patrick Valduriez.

While structured systems are well-studied in the P2P research community, unstructured systems are still the most deployed systems in today's internet. Unstructured systems are attractive because of their simplicity and their high robustness. However, they suffer from high bandwidth consumption which remains a salient issue for both users and internet providers. In fact, in unstructured networks, there is no control on either the overlay topology or the data placement, and thus queries are non-deterministically routed. Basically, the fundamental routing mechanisms are flooding and its variations (e.g. random walks). The major concern with these mechanisms is the total number of messages generated per query. More severely, a large fraction of these messages are redundant, and unnecessarily increase the download on the system. This is due to the adhoc nature of P2P connections and the lack of structural information about the network topology. In our work, we mainly aim to improve the performance of flooding-based mechanisms, by simply exploiting the characteristics of the overlay topology, and this irrespective of the use of data indexes. Our approach consists in eliminating the redundancy in query messages while keeping a good network coverage, i.e. reaching a large number of nodes, which is an interesting merit of flooding approaches.

To this end, we proposed CBST [63], a Cluster-Based Search Technique for P2P systems, which relies on connectivity-based clustering schemes. Such clustering schemes are considered as a way to extract inherent structural patterns from the P2P network overlay. The network is partitioned into clusters based on node connectivity, such that nodes within clusters are highly connected, while nodes between clusters are loosely connected. The CBST technique works as follows. Based on a local knowledge about the network clustering in its neighborhood, each node maintains a spanning tree (rooted at itself) over its own cluster, and information on global links through which it can reach other clusters with minimal costs. These information are efficiently gathered and maintained in a cluster-based routing table. The benefits in query routing are two folds. First, a query is efficiently disseminated in a cluster, without any redundant messages, using the spanning tree of the first node contacted in that cluster. Second, the query messages between clusters are restricted to those traversing the global links provided by the routing table of the querying node. Extensive simulations have demonstrated the efficiency of the CBST technique compared to the pure flooding and random walk routing techniques.

### *6.3.3. Testing P2P Systems*

**Participants:** Eduardo Almadeia, Gerson Sunyé, Patrick Valduriez.

Testing P2P systems is hard because they must be deployed on a high number of nodes, which can be autonomous, refusing to answer to some requests or even unexpectedly leaving the system. Such volatility of nodes is a common behavior in P2P systems and can be interpreted as fault during tests. To address P2P testing, we proposed algorithms to execute test cases, test architectures and a new methodology. This work is done in cooperation with Yves Le Traon (Triskell project-team).

In [49], we proposed a synchronization algorithm for executing test case actions in P2P systems. The main goal of the algorithm is to progressively dispatch the actions of a test case to a set of nodes and ensure that all nodes complete the execution of an action, even upon volatility, before dispatching the next one. In [48], we proposed a framework for testing P2P systems that can combine the execution of functional tests to the simulation of volatility. The capabilities of this framework are (1) to automate the execution of each local-to-a-peer test case, (2) to build automatically the global verdict, (3) to allow the explicit control of each peer volatility.

In [50], we proposed two different architectures to synchronize the execution of test cases in large-scale P2P systems. The first architecture organizes the testers in a balanced tree (B-Tree) structure where the synchronization is performed from the root to the leaves. The second approach uses gossiping messages among testers, reducing communication among the testers responsible to execute consecutive test case actions. Since both architectures do not rely on a central coordinator, they scale up well.

In [47], we proposed an incremental methodology to deal with three aspects of P2P testing (functionality, volatility and scalability). The idea is to cover functionality first on a small system and then incrementally address the scalability and volatility aspects. Empirical results obtained by running four test cases on a popular P2P system illustrate the fact that satisfying a simple test criterion such as code coverage is a hard task. We also identified open issues, such as the generation of efficient test objectives. We validated our framework through implementation and experimentation on two popular open-source P2P systems (FreePastry and OpenChord). Through experimentation, we analyzed the behavior of both systems on different conditions of volatility and showed how the framework is able to detect implementation problems and cope with a large-scale P2P system under test.

## 6.4. Data Access using DHTs

An important kind of P2P systems is DHTs which provide an efficient hash table interface to return contents (e.f. a file) for a given key (e.g. file name). Queries can be efficiently routed since the routing scheme allows one to find a peer responsible for a key in $O(logn)$ routing hops, where $n$ is the number of peers in the network. Thus, a DHT can be viewed as a fast, decentralized index and is useful to support efficiently various data management fucntions. We exploit DHTs for improving the performance of data streaming applications and to improve data privacy of P2P systems.

### *6.4.1. Data streaming*

**Participants:** Esther Pacitti, Wenceslao Palma, Patrick Valduriez.

Many important applications such as network monitoring, sensor networks, financial applications generate streams of data. Processing a query over a data stream involves running the query continuously over the stream and generating a new answer each time a new data item arrives. Due to the unbounded nature of data streams, it is not possible to store the data entirely in a bounded memory. This makes difficult the processing of queries that need to compare each new arriving data with past ones. As an example, consider a network monitoring application that needs to issue a join query over traffic traces from various links, in order to monitor the total traffic that is common among three links L1, L2 and L3 over the last 10 minutes.

A common solution to the problem of processing join queries over data streams is to execute the query over a sliding window that maintains a restricted number of recent data items. This allows queries to be executed in a finite memory and in an incremental manner by generating new answers when a new data item arrives. In [35], we addressed the problem of computing approximate answers to windowed stream joins over data streams. Our solution involves a scalable distributed sliding window that takes advantage of the indexing

power of DHT networks and can be equivalent to thousands of centralized sliding windows. We proposed a method, called DHTJoin, which deals with efficient execution of join queries over all data items which are stored in the distributed sliding window. DHTJoin combines hash-based placement of tuples in the DHT and dissemination of queries using a gossip style protocol. We evaluated the performance of DHTJoin through simulation. The results show the effectiveness of our solution compared with previous work. This work was done in cooperation with R. Akbarinia (Univ. Waterloo).

### *6.4.2. Data Privacy*

**Participants:** Mohamed Jawad, Patrick Valduriez.

In a P2P system, when sharing data for different *purposes* (e.g., billing, purchase, shipping, etc.), data privacy can be easily violated by untrustworthy peers wich may use data for other purposes (e.g., marketing, fraudulence, profiling, etc.). A basic principle of data privacy is *purpose* specification which states that data providers should be able to specify the purpose for which their data will be collected and used. In the context of P2P systems, decentralized control makes it hard to enforce purpose-based privacy. And the major problem of data disclosure is not addressed. Hippocratic databases provide mechanisms for enforcing *purpose-based* disclosure control, within a centralized datastore. This is achieved by using privacy metadata, i.e. privacy policies and privacy authorizations stored in tables. A privacy policy defines for each attribute, tuple or table the usage purpose(s), the potential users and retention period while privacy authorization defines which purposes each user is authorized to use.

In [33], we apply the Hippocratic database principles to P2P systems to enforce purpose-based privacy. We focus on Distributed Hash Tables (DHTs), because they provide strong guarantees in terms of access performance. More precisely, we exploit DHTs to store only privacy metadata, so that purpose-based control can be decentralized and efficient. However, private user data remain in user peers which can maintain strong local control. This allows for instance to use highly secure solutions locally such as secure-chip DBMS (see results of recent joint work with the Smis project-team in [16]).

We also proposed PriServ, a privacy service which prevents privacy violation by prohibiting malicious data access, using purpose-based access control and trust among peers. The performance evaluation of our approach through simulation shows that the overhead introduced by PriServ is small.

# 7. Contracts and Grants with Industry

## 7.1. STREP Grid4All (2006-2009)

**Participants:** William Kokou Dedzoe, Philippe Lamarre, Esther Pacitti, Jorge Quiane, Patrick Valduriez.

The project is with France Telecom R&D (leader), INRIA (Atlas, Grand-Large, Regal, and Sardes), Kungliga Tekniska Hoegskolan, Swedish Institute of Computer Science, ICCS (Greece), University of Piraeus Research Center, Universitat Politecnica 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).

## 7.2. RNTL XWiki Concerto (2006-2008)

**Participants:** 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.

## 7.3. 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. Other Grants and Activities

## 8.1. Regional Actions

We are involved in the following actions:

### *8.1.1. MILES (2007–2010)*

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 above) is further supported by Pôle de Compétitivité Images & Réseaux.

## 8.2. National Actions

We are involved in the following projects:

### *8.2.1. ARA Massive Data Respire (2006-2008)*

**Participants:** Philippe Lamarre, Esther Pacitti, Jorge Quiane, Gerson Sunyé, 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. ARA Massive Data SemWeb (2004-2008)*

**Participant:** 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.

### *8.2.3. ARA Massive Data APMD (2004-2008)*

**Participant:** Guillaume Raschia.

The project APMD (Personalised Access to 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 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 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.

Furthermore, we have regular scientific relationships with research laboratories in

- North America: Univ. of Waterloo (Reza Akbarinia, Tamer Özsu), McGill University (Bettina Kemme);
- Europe: Univ. of Madrid (Ricardo Jimenez-Periz), Univ. of Barcelona (Josep Lluis Larriba Pey), Univ. of Roskilde (Henrik Larsen), Nokia (Andreas Myka), ;
- Brazil: Univ. Federal of Rio de Janeiro (Marta Mattoso), PUC-Rio (Sergio Lifschitz), PUCPR, Curitiba, Brazil (Vidal Martins);
- International Univ. of Rabat (Noureddine Mouaddib, formely in the team).

# 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 BDA.

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

## 9.2. Editorial Program committees

Participation in the editorial board of scientific journals:

- Proceedings of the VLDB Endowment: P. Valduriez.
- Distributed and Parallel Database Systems, Kluwer Academic Publishers: P. Valduriez.
- Internet and Databases: Web Information Systems, Kluwer Academic Publishers: P. Valduriez.
- Book series “Data Centric Systems and Applications” (Springer-Verlag): P. Valduriez.
- Ingénierie des Systèmes d’Information, Hermès : P. Valduriez.

Participation in conference program committees :

- ACM-SIGMOD Int. Conf. 2008: P. Valduriez.
- Int. Conf. on VLDB 2008: P. Valduriez; 2009: E. Pacitti.
- ACM Int. Conf. on Information and Knowledge Management (CIKM) 2008: E. Pacitti.
- European Dependable Computing Conference (EDCC) 2008: E. Pacitti.
- IEEE Int. Conf. on Distributed Computing Systems (ICDCS) 2009, Data Management track: E. Pacitti.
- Int. Conf. on High Performance Computing for Computational Science (VecPar) 2008: P. Valduriez.

- Brazilian Symposium on Databases (SBBD) 2008: E. Pacitti.
- Damap workshop on P2P data management, co-located with EDBT 2008: E. Pacitti (PC co-chair), P. Lamarre.
- 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).
- European Conf. on Parallel Computing (Euro-Par) 2008, Parallel and Distributed Database Track: E. Pacitti (vice chair).
- Int. Conf. on Extending DataBase Technologies (EDBT) 2008: E. Pacitti, G. Raschia; 2009: E. Pacitti.
- Int. Conf. on Advanced Information Systems Engineering (CAiSE) 2008: P. Valduriez.
- IEEE Journée Francophone sur la Cohérence de Données en Univers Reparti (CDUR), 2008: E. Pacitti.
- Journées Bases de Données Avancées (BDA), 2008: P. Lamarre.
- Artificial Economics (AE), 2008, P. Lamarre.

## 9.3. Invited Talks

In April, E. Pacitti and P. Valduriez visited the university of Barcelona (Prof. J.L. Larriba Pey ), for a week and gave talks on query processing and replication in P2P.

In June, E. Pacitti gave a keynote talk on Grid Data Management at the Int. Workshop on High-Performance Data Management in Grid Environments (HPDGrid).

In July, P. Valduriez gave invited talks on the DataRIng project at UFRJ, Rio de Janeiro and Univ. of Recife, Pernanbuco, Brazil.

## 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. Noureddine Mouaddib, now on leave from Polytech'Nantes and a former member of the team, is starting a new international university in Rabbah, Moroco, with the objectives of training high-level managers for Africa.

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. A third edition is in progress and will be a major revision with much new material on replication, P2P, parallel systems and web data integration.

# 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] A. NIKSERESHT, M. GELGON. *Gossip-based Computation of a Gaussian Mixture Model for Distributed Multimedia Indexing*, in "IEEE Transactions on Multimedia", vol. 10, n$^{o}$ 3, April 2008, p. 385-392.

[5] 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.

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

[7] J.-A. 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.

[8] J.-A. QUIANÉ-RUIZ, P. LAMARRE, P. VALDURIEZ. *A Self-Adaptable Query Allocation Framework for Distributed Information Systems*, in "The VLDB Journal", 2008.

[9] 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.

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

# Year Publications

## Doctoral Dissertations and Habilitation Theses

[11] A. NIKSERESHT. *Estimation de modèles de mélanges probabilistes : une proposition pour un fonctionnement réparti et décentralisé*, Ph. D. Thesis, Université de Nantes, 2008.

[12] E. PACITTI. *Réplication asynchrone des données dans trois contextes: entrepôts, grappes et systèmes pair-à-pair*, Habilitation Thesis, Ph. D. Thesis, Université de Nantes, 2008.

[13] J.-A. QUIANÉ-RUIZ. *Allocation de requêtes dans les systèmes d'information distribués autonomes*, Ph. D. Thesis, Université de Nantes, 2008.

[14] J. ROUGUI. *Indexation des documents audio : traitement des difficultés apparaissant avec les grands volumes de données*, Ph. D. Thesis, Université de Nantes, 2008.

[15] A. VENTRESQUE. *Espaces vectoriels sémantiques: enrichissement et interprétation de requêtes dans un système d'informations distribué hétérogène*, Ph. D. Thesis, Université de Nantes, 2008.

## Articles in International Peer-Reviewed Journal

[16] N. ANCIAUX, L. BOUGANIM, P. PUCHERAL, P. VALDURIEZ. *DiSC: Benchmarking Secure Chip DBMS*, in "IEEE Trans. on Knowledge and Data Engineering", vol. 20, n$^{o}$ 10, 2008, p. 1363-1377.

[17] M. DIDONET DEL FABRO, P. VALDURIEZ. *Towards the Efficient Development of Model Transformations using Model Weaving and Matching Transformations*, in "Software and Systems Modeling (SoSyM)", 2008, http://www.springerlink.com/content/n700668617702610/.

[18] 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° 4, 2008, p. 251-262.

[19] R. HAYEK, G. RASCHIA, P. VALDURIEZ, N. MOUADDIB. *Towards a Data Summarization Service in APPA*, in "Int. Journal of Pervasive Computing and Communications (IJPCC)", to appear, vol. Special Issue on Towards Merging Grid and Pervasive Computing, 2008.

[20] N. KOTOWSKI, A. LIMA, E. PACITTI, P. VALDURIEZ, M. MATTOSO. *OLAP Query Processing in Grids*, in "Concurrency and Computation: Practice and Experience", Selection of Best Papers of the VLDB Data Management in Grids Workshop (VLDB DMG 2007), vol. 20, n° 17, 2008, p. 2039-2048.

[21] V. MARTINS, P. PACITTI, M. EL DICK, R. JIMÉNEZ-PERIS. *Scalable and Topology-Aware Reconciliation in P2P Networks*, in "Distributed and Parallel Databases (DAPD)", vol. 24, n° 1-3, 2008, p. 1-43.

[22] A. NIKSERESHT, M. GELGON. *Gossip-based computation of a Gaussian mixture model for distributed multimedia indexing*, in "IEEE Transactions on Multimedia", vol. 10, n° 3, 2008, p. 385-392.

[23] J.-A. QUIANÉ-RUIZ, P. LAMARRE, P. VALDURIEZ. *A Self-Adaptable Query Allocation Framework for Distributed Information Systems*, in "The VLDB Journal", 2008.

[24] L. UGHETTO, A. VOGLOZIN, N. MOUADDIB. *Database querying with personalized vocabulary using data summaries*, in "Int. Journal of Fuzzy Sets and Systems", vol. 159, n° 15, 2008, p. 2030-2048.

### Articles in National Peer-Reviewed Journal

[25] R. HAYEK, G. RASCHIA, P. VALDURIEZ, N. MOUADDIB. *Gestion de résumés de données dans les systèmes pair-à-pair non structurés*, in "Ingénierie des Systèmes d'Information (ISI)", to appear, vol. Numéro spécial Networking and Information Systems, 2008.

### International Peer-Reviewed Conference/Proceedings

[26] P. BRUNEAU, M. GELGON, F. PICAROUGNE. *Parameter-based reduction of Gaussian mixture models with a variational-Bayes approach*, in "Int. Conf. on Pattern Recognition (ICPR'08), Tampa (Fl.), USA", to appear, 2008.

[27] P. BRUNEAU, A. PIGEAU, M. GELGON, F. PICAROUGNE. *Geo-temporal structuring of a personal image database with two-level variational-Bayes mixture estimation*, in "Adaptive Multimedia Retrieval workshop (AMR'08), to appear as LNCS post-proceedings, Berlin, Germany", 2008.

[28] G. CANALS, P. MOLLI, J. MAIRE, S. LAURIÈRE, E. PACITTI, M. TLILI. *XWiki Concerto: A P2P Wiki System Supporting Disconnected Work*, in "Int. Conf. on Cooperative Design, Visualization, and Engineering (CDVE 2008), Calvià, Mallorca, Spain", LNCS, vol. 5220, Springer, 2008, p. 98-106.

[29] M. EL DICK, E. PACITTI, B. KEMME. *Flower-CDN: a Hybrid P2P Overlay for Efficient Query Processing in CDN*, in "Int. Conf. on Extending Database Technology (EDBT), Saint Pertersburg, Russia", ACM International Conference Proceeding Series, To appear, 2009.

[30] 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, 2008.

[31] 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", 2008, p. 422-428.

[32] R. HAYEK, G. RASCHIA, P. VALDURIEZ, N. MOUADDIB. *Summary Management in P2P Systems*, in "Int. Conf. on Extending Database Technology (EDBT), Nantes, France", ACM International Conference Proceeding Series, 2008, p. 16-25.

[33] M. JAWAD, P. SERRANO-ALVARADO, P. VALDURIEZ. *Design of PriServ, a privacy service for DHTs*, in "Int. Workshop on Privacy and Anonymity in Information Society (PAIS), Nantes", ACM International Conference Proceeding Series, 2008, p. 21-25.

[34] M. PAES, A. LIMA, P. VALDURIEZ, M. MATTOSO. *High-performance Query Processing of a Real-world OLAP Database with ParGRES*, in "Int. Conf. on High Performance Computing for Computational Science (VecPar), Toulouse", LNCS, Best student paper award, vol. 5336, Springer, 2008, p. 192-205.

[35] W. PALMA, R. AKBARINIA, E. PACITTI, P. VALDURIEZ. *Efficient Processing of Continuous Join Queries Using Distributed Hash Tables*, in "European Conf. on Parallel Computing (Euro-Par)", 2008, p. 632-641.

[36] K. PARK, P. VALDURIEZ, H. CHOO. *Mobile Continuous Nearest Neighbor Queries on Air*, in "ACM SIGSPATIAL Int. Conf. on Advances in Geographic Information Systems (GIS), Irvine, California", 2008.

[37] Q.-K. PHAM, R. SAINT-PAUL, B. BENATALLAH, N. MOUADDIB, G. RASCHIA. *Mine your own business, mine others' news!*, in "Int. Conf. on Extending Database Technology (EDBT), Nantes, France", ACM, 2008, p. 725-729.

[38] J.-A. QUIANÉ-RUIZ, P. LAMARRE, S. CAZALENS, P. VALDURIEZ. *Managing Virtual Money for Satisfaction and Scale up in P2P Systems*, in "Int. Workshop on Data Management in Peer-to-Peer Systems (DaMaP), Nantes", ACM International Conference Proceeding Series, 2008, p. 67-74.

[39] J.-A. QUIANÉ-RUIZ, P. LAMARRE, P. VALDURIEZ. *SbQA: A Self-Adaptable Query Allocation Process*, in "IEEE Int. Conf. on Data Engineering (ICDE'09)", to appear (Demo paper), 2008.

[40] J.-A. QUIANÉ-RUIZ, P. LAMARRE, P. VALDURIEZ. *SbQA: Une méthode auto-adaptative pour l'allocation de requêtes*, in "Journées Bases de Données Avancées (BDA)", Demo paper, 2008.

[41] J. R. M. SANCHEZ, J. MARTINEZ, P. VALDURIEZ. *Efficient Processing of Nearest Neighbor Queries in Parallel Multimedia Databases*, in "Int. Conf. on Database and Expert Systems Applications (DEXA), Turin, Italy", LNCS, vol. 5181, Springer, 2008, p. 326-339.

[42] M. TLILI, W. K. DEDZOE, E. PACITTI, R. AKBARINIA, P. VALDURIEZ, P. MOLLI, G. CANALS, S. LAURIÈRE. *Estampillage et Journalisation P2P pour Xwiki*, in "Conférence Internationale sur les NOuvelles TEchnologies de la REpartition (NOTERE), Lyon", ACM, 2008, p. 197-200.

[43] M. TLILI, W. K. DEDZOE, E. PACITTI, R. AKBARINIA, P. VALDURIEZ, P. MOLLI, G. CANALS, S. LAURIÈRE. *P2P Logging and Timestamping for Reconciliation*, in "Proceedings of the VLDB Endowment (Papers from VLDB2008 Conf.), Auckland, New Zeland", vol. 1, n$^{o}$ 2, 2008, p. 1420-1423.

[44] A. VENTRESQUE, S. CAZALENS, P. LAMARRE, P. VALDURIEZ. *Dealing with P2P Semantic Heterogeneity through Query Expansion and Interpretation*, in "Int. Workshop on Data Management in Peer-to-Peer Systems (DaMaP), Nantes", ACM International Conference Proceeding Series, 2008, p. 3-10.

[45] A. VENTRESQUE, S. CAZALENS, P. LAMARRE, P. VALDURIEZ. *Improving Interoperability Using Query Interpretation in Semantic Vector Spaces*, in "European Semantic Web Conference (ESWC), Tenerife, Canary Islands, Spain", LNCS, Nominated for Best Paper Award (top 4), vol. 5021, Springer, 2008, p. 539-553.

[46] A. VENTRESQUE, S. CAZALENS, P. LAMARRE, P. VALDURIEZ. *Interprétation de requêtes pour aider les pairs à se comprendre dans des systèmes sémantiquement hétérogènes*, in "Journés Bases de Données Avancées (BDA), Guilherand-Granges, France", 2008.

[47] E. C. DE ALMEIDA, G. SUNYÉ, Y. L. TRAON, P. VALDURIEZ. *A Framework for Testing Peer-to-Peer Systems*, in "Int. Symp. on Software Reliability Engineering (ISSRE), Redmond, Seattle", 2008.

[48] E. C. DE ALMEIDA, G. SUNYÉ, Y. L. TRAON, P. VALDURIEZ. *Testing Peers' Volatility*, in "IEEE/ACM Int. Conf. on Automated Software Engineering (ASE), L'Aquila, Italy", 2008.

[49] E. C. DE ALMEIDA, G. SUNYÉ, P. VALDURIEZ. *Action Synchronization in P2P System Testing*, in "Int. Workshop on Data Management in Peer-to-Peer Systems (DaMaP), Nantes", ACM International Conference Proceeding Series, 2008, p. 43-49.

[50] E. C. DE ALMEIDA, G. SUNYÉ, P. VALDURIEZ. *Testing Architectures for Large Scale Grids*, in "Int. Workshop on High-Performance Data Management in Grid Environments (HPDGrid), Toulouse", LNCS, vol. 5336, Springer, 2008, p. 555-566.

[51] D. DE SOUSA, S. LIFSCHITZ, P. VALDURIEZ. *BLAST Distributed Execution on Partitioned Databases with Primary Fragments*, in "Int. Workshop on High-Performance Data Management in Grid Environments (HPDGrid), Toulouse", LNCS, vol. 5336, Springer, 2008, p. 544-554.

### National Peer-Reviewed Conference/Proceedings

[52] A. VENTRESQUE, S. CAZALENS, P. LAMARRE, P. VALDURIEZ. *Enrichissement sémantique de requête utilisant un ordre sur les concepts*, in "Atelier "Mesures de similarité sémantique", associé à la conférence Extraction et Gestion des Connaissances (EGC), Sophia-Antipolis, France", 2008.

### Workshops without Proceedings

[53] P. BRUNEAU, F. PICAROUGNE, M. GELGON. *Incremental clustering in a data stream with a flock of agents*, in "Proc. of joint meeting of the Société Francophone de Classification and the Classification and Data Analysis Group of the Italian Society of Statistics (SFC-CLADAG), Caserta, Italy", June 2008, p. 192-198.

### Scientific Books (or Scientific Book chapters)

[54] R. HAYEK, G. RASCHIA, P. VALDURIEZ, N. MOUADDIB. *Data Localization and Description Through Summaries in P2P Collaborative Applications*, in "Handbook of Peer-to-Peer Networking", X. (. SHEN, H. YU, J. BUFORD, M. AKON (editors), to appear, Springer, 2008.

[55] N. MOUADDIB, G. RASCHIA, A. VOGLOZIN, L. UGHETTO. *From User Requirements to Evaluation Strategies of Flexible Queries in Databases*, in "Handbook of Research on Fuzzy Information Processing in Databases", J. GALINDO (editor), n° 5, IGI Global, 2008, p. 115-142.

[56] E. PACITTI. *Parallel Query Processing*, in "Encyclopedia of Database Systems", L. LIU, M. ÖZSU (editors), to appear, Springer, 2008.

[57] P. VALDURIEZ. *Parallel Data Placement*, in "Encyclopedia of Database Systems", L. LIU, M. ÖZSU (editors), to appear, Springer, 2008.

[58] P. VALDURIEZ. *Parallel Database Management*, in "Encyclopedia of Database Systems", L. LIU, M. ÖZSU (editors), to appear, Springer, 2008.

### Books or Proceedings Editing

[59] A. DOUCET, S. GANÇARSKI, E. PACITTI (editors). *Proceedings of the 2008 International Workshop on Data Management in Peer-to-Peer Systems, DaMaP 2008, Nantes, France, March 25, 2008*, ACM International Conference Proceeding Series, ACM, 2008.

### Research Reports

[60] M. BECHCHI, G. RASCHIA, N. MOUADDIB. *Joining Distribued Database Summaries*, Technical report, n° 6768, INRIA, 2008, http://hal.inria.fr/inria-00346528/fr/.

[61] M. BECHCHI, A. VOGLOZIN, G. RASCHIA, N. MOUADDIB. *Multi-dimensional Grid-based Clustering of Fuzzy Query Results*, Technical report, n° 6770, INRIA, 2008, http://hal.inria.fr/inria-00346540/fr/.

[62] M. EL DICK, E. PACITTI, B. KEMME. *Flower-CDN: a Hybrid P2P Overlay for Efficient Query Processing in CDN*, Technical report, n° 6689, INRIA, 2008, http://hal.inria.fr/inria-00331231/fr/.

[63] R. HAYEK, G. RASCHIA, P. VALDURIEZ. *Cluster-based Search Technique for P2P Systems*, Technical report, n° 6782, INRIA, 2008, http://hal.inria.fr/inria-00348313/fr/.

### Other Publications

[64] A. MAILLET, F. RIVARD, F. AUTRUSSEAU, R. LEHN, M. GELGON, P. BRUNEAU, F. PICAROUGNE. *The Safimage Platform*, Demo paper, October 2008, Network and Electronic Media Forum (NEM), Saint-Malo, France.