IN PARTNERSHIP WITH:
Université des sciences et techniques du Languedoc (Montpellier 2)

Activity Report 2011

Team ZENITH

Scientific Data Management

RESEARCH CENTER
Sophia Antipolis - Méditerranée

THEME
Knowledge and Data Representation and Management

Table of contents

Team ZENITH

Keywords: Data Management, Scientific Data, Information Indexing And Retrieval, Workflow, Parallelism

1. Members

Research Scientists

Reza Akbarinia [Junior Researcher, INRIA]
Alexis Joly [Junior Researcher, INRIA, joined Zenith in July]
Florent Masseglia [Junior Researcher, INRIA]
Didier Parigot [Junior Researcher, INRIA]
Patrick Valduriez [Team Leader, Senior Researcher, INRIA, HdR]

Faculty Members

Zohra Bellahsène [Professor, UM2, HdR]
Hinde Bouziane [Associate Professor, UM2]
Michèle Cart [Associate Professor, UM2]
Rémi Coletta [Associate Professor, UM2 (INRIA delegation until June)]
Esther Pacitti [Associate Team Leader, Professor, UM2, HdR]

Technical Staff

Emmanuel Castanier [Engineer, ADT WebSmatch]
Julien Champ [Engineer, ANR Otmedia, since September]
Guillaume Verger [Engineer, ANR DataRing]

PhD Students

Ayoub Ait Lahcen [ANR STAMP fellowship]
William Kokou Dedzoe [CNRS fellowship, LINA]
Fady Draidi [MAE-INRIA fellowship]
Vincenzo Gulisano [Universidad Politecnica de Madrid, Spain]
Mohamed Jawad [University of Nantes fellowship, LINA]
Pierre Letessier [CIFRE fellowship, Telecom ParisTech]
Miguel Liroz [INRIA CORDIS fellowship]
Imen Mami [UM2 fellowship]
Duy Hoa Ngo [ANR DataRing fellowship]
Eduardo Ogasawara [Universidade Federal de Rio de Janerio, Brazil]
Toufik Sarni [CNRS fellowship, LINA]
Maximilien Servajean [NUMEV-UM2 fellowship]
Mounir Tlili [Pays-de-la-Loire fellowship, LINA]

Administrative Assistant

Annie Aliaga

2. Overall Objectives

2.1. Introduction

Modern science such as agronomy, bio-informatics, astronomy and environmental science must deal with overwhelming amounts of experimental data produced through empirical observation and simulation (http://www.computational-sustainability.org). Such data must be processed (cleaned, transformed, analyzed) in all kinds of ways in order to draw new conclusions, prove scientific theories and produce knowledge. However, constant progress in scientific observational instruments (e.g. satellites, sensors, large hadron collider) and simulation tools (that foster in silico experimentation, as opposed to traditional in situ or in vivo experimentation) creates a huge data overload. For example, climate modeling data are growing so fast that they will lead to collections of hundreds of exabytes (10^{18} bytes) expected by 2020.

Scientific data is also very complex, in particular because of heterogeneous methods used for producing data, the uncertainty of captured data, the inherently multi-scale nature (spatial scale, temporal scale) of many sciences and the growing use of imaging (e.g. satellite images), resulting in data with hundreds of attributes, dimensions or descriptors. Processing and analyzing such massive sets of complex scientific data is therefore a major challenge since solutions must combine new data management techniques with large-scale parallelism in cluster, grid or cloud environments.

Furthermore, modern science research is a highly collaborative process, involving scientists from different disciplines (e.g. biologists, soil scientists, and geologists working on an environmental project), in some cases from different organizations distributed in different countries. Since each discipline or organization tends to produce and manage its own data, in specific formats, with its own processes, integrating distributed data and processes gets difficult as the amounts of heterogeneous data grow.

Despite their variety, we can identify common features of scientific data: massive scale; manipulated through complex, distributed workflows; typically complex, e.g. multidimensional or graph-based; with uncertainty in the data values, e.g., to reflect data capture or observation; important metadata about experiments and their provenance; and mostly append-only (with rare updates).

Generic data management solutions (e.g. relational DBMS) which have proved effective in many application domains (e.g. business transactions) are not efficient for dealing with scientific data, thereby forcing scientists to build ad-hoc solutions which are labor-intensive and cannot scale. In particular, relational DBMSs have been lately criticized for their “one size fits all” approach. Although they have been able to integrate support for all kinds of data (e.g., multimedia objects, XML documents and new functions), this has resulted in a loss of performance and flexibility for applications with specific requirements because they provide both “too much” and “too little”. Therefore, it has been argued that more specialized DBMS engines are needed. For instance, column-oriented DBMSs, which store column data together rather than rows in traditional row-oriented relational DBMSs, have been shown to perform more than an order of magnitude better on decision-support workloads. The “one size does not fit all” counter-argument generally applies to cloud data management as well. Cloud data can be very large, unstructured (e.g. text-based) or semi-structured, and typically append-only (with rare updates). And cloud users and application developers may be in high numbers, but not DBMS experts. Therefore, current cloud data management solutions have traded consistency for scalability, simplicity and flexibility. As alternative to relational DBMS (which use the standard SQL language), these alternative solutions have been quoted as Not Only SQL (NOSQL) by the database research community.

The three main challenges of scientific data management can be summarized by: (1) scale (big data, big applications); (2) complexity (uncertain, multi-scale data with lots of dimensions), (3) heterogeneity (in particular, data semantics heterogeneity). The overall goal of Zenith is to address these challenges, by proposing innovative solutions with significant advantages in terms of scalability, functionality, ease of use, and performance. To produce generic results, these solutions will be in terms of architectures, models and algorithms that can be implemented in terms of components or services in specific computing environments, e.g. grid, cloud. To maximize impact, a good balance between conceptual aspects (e.g. algorithms) and practical aspects (e.g. software development) is necessary. We plan to design and validate our solutions by working closely with scientific application partners (CIRAD, INRA, CEMAGREF, etc.). To further validate our solutions and extend the scope of our results, we also want to foster industrial collaborations, even in non scientific applications, provided that they exhibit similar challenges.

2.2. Highlights

The third edition of the Özsu-Valduriez textbook *Principles of Distributed Database Systems* [11] has been released (by Springer). This long awaited major revision is now about 850 pages. In addition to the fundamental principles of distributed data management, it now covers new hot topics such as web data management, peer-to-peer, data streaming, and cloud.

At the 2011 competition of the Ontology Alignment Evaluation Initiative (http://oaei.ontologymatching.org), our YAM++ ontology matching tool achieved excellent results: first position at the Conference track and second position at the Benchmark track.

The paper [46] on aggregate queries in uncertain databases was selected as one of the best papers at BDA 2011.

3. Scientific Foundations

3.1. Data Management

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, search engines, 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 modeling, 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 (smart phone, 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.

The problems of scientific data management (massive scale, complexity and heterogeneity) go well beyond the traditional context of DBMS. To address them, we capitalize on scientific foundations in closely related domains: distributed data management, cloud data management, uncertain data management, metadata integration, data mining and content-based information retrieval.

3.2. Distributed Data Management

To deal with the massive scale of scientific data, we exploit large-scale distributed systems, with the objective of making distribution transparent to the users and applications. Thus, we capitalize on the principles of large-scale distributed systems such as clusters, peer-to-peer (P2P) and cloud, to address issues in data integration, scientific workflows, recommandation, query processing and data analysis, .

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) [11]. 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, e.g. price comparators such as KelKoo, extend the distributed database approach to access data sources on the Internet with a simpler query language in read-only mode.

Parallel database systems 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.

Scientific workflow management systems (SWfMS) such as Kepler (http://kepler-project.org) and Taverna (http://www.taverna.org.uk) allow scientists to describe and execute complex scientific procedures and activities, by automating data derivation processes, and supporting various functions such as provenance management, queries, reuse, etc. Some workflow activities may access or produce huge amounts of distributed data and demand high performance computing (HPC) environments with highly distributed data sources and computing resources. However, combining SWfMS with HPC to improve throughput and performance remains a difficult challenge. In particular, existing workflow development and computing environments have limited support for data parallelism patterns. Such limitation makes complex the automation and ability to perform efficient parallel execution on large sets of data, which may significantly slow down the execution of a workflow.

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 BitTorrent 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, whereby peers forward messages to their neighbors. This work led to structured solutions based on Distributed Hash Tables (DHT), e.g. CHORD and Pastry, or hybrid solutions with super-peers that index subsets of peers. Another approach is to exploit gossiping protocols, also known as epidemic protocols. Gossiping has been initially proposed to maintain the mutual consistency of replicated data by spreading replica updates to all nodes over the network. It has since been successfully used in P2P networks for data dissemination. Basic gossiping is simple. Each peer has a complete view of the network (i.e. a list of all peers' addresses) and chooses a node at random to spread the request. The main advantage of gossiping is robustness over node failures since, with very high probability, the request is eventually propagated to all nodes in the network. In large P2P networks, however, the basic gossiping model does not scale as maintaining the complete view of the network at each node would generate very heavy communication traffic. A solution to scalable gossiping is by having each peer with only a partial view of the network, e.g. a list of tens of neighbor peers. To gossip a request, a peer chooses at random a peer in its partial view to send it the request. In addition, the peers involved in a gossip exchange their partial views to reflect network changes in their own views. Thus, by continuously refreshing their partial views, nodes can self-organize into randomized overlays which scale up very well.

We claim that a P2P solution is the right solution to support the collaborative nature of scientific applications as it provides scalability, dynamicity, autonomy and decentralized control. Peers can be the participants or organizations involved in collaboration and may share data and applications while keeping full control over their (local) data sources.

But for very-large scale scientific data analysis or to execute very large data-intensive workflow activities (activities that manipulate huge amounts of data), we believe cloud computing (see next section), is the right approach as it can provide virtually infinite computing, storage and networking resources. However, current cloud architectures are proprietary, ad-hoc, and may deprive users of the control of their own data. Thus, we postulate that a hybrid P2P/cloud architecture is more appropriate for scientific data management, by combining the bests of both. In particular, it will enable the clean integration of the users' own computational resources with different clouds.

3.3. Cloud Data Management

Cloud computing encompasses on demand, reliable services provided over the Internet (typically represented as a cloud) with easy access to virtually infinite computing, storage and networking resources. Through very simple Web interfaces and at small incremental cost, users can outsource complex tasks, such as data storage, system administration, or application deployment, to very large data centers operated by cloud providers. Thus, the complexity of managing the software/hardware infrastructure gets shifted from the users' organization to the cloud provider. From a technical point of view, the grand challenge is to support in a cost-effective way the very large scale of the infrastructure which has to manage lots of users and resources with high quality of service.

Cloud customers could move all or part of their information technology (IT) services to the cloud, with the following main benefits:

- **Cost.** The cost for the customer can be greatly reduced since the IT infrastructure does not need to be owned and managed; billing is only based only on resource consumption. For the cloud provider, using a consolidated infrastructure and sharing costs for multiple customers reduces the cost of ownership and operation.
- **Ease of access and use.** The cloud hides the complexity of the IT infrastructure and makes location and distribution transparent. Thus, customers can have access to IT services anytime, and from anywhere with an Internet connection.
- **Quality of Service (QoS).** The operation of the IT infrastructure by a specialized provider that has extensive experience in running very large infrastructures (including its own infrastructure) increases QoS.
- **Elasticity.** The ability to scale resources out, up and down dynamically to accommodate changing conditions is a major advantage. In particular, it makes it easy for customers to deal with sudden increases in loads by simply creating more virtual machines.

However, cloud computing has some drawbacks and not all applications are good candidates for being "cloudified". The major concern is wrt. data security and privacy, and trust in the provider (which may use no so trustful providers to operate). One earlier criticism of cloud computing was that customers get locked in proprietary clouds. It is true that most clouds are proprietary and there are no standards for cloud interoperability. But this is changing with open source cloud software such as Hadoop, an Apache project implementing Googgle's major cloud services such as Google File System and MapReduce, and Eucalyptus, an open source cloud software infrastructure, which are attracting much interest from research and industry.

There is much more variety in cloud data than in scientific data since there are many different kinds of customers (individuals, SME, large corporations, etc.). However, we can identify common features. Cloud data can be very large, unstructured (e.g. text-based) or semi-structured, and typically append-only (with rare updates). And cloud users and application developers may be in high numbers, but not DBMS experts.

Current cloud data management (NOSQL) solutions typically trade consistency for scalability, simplicity and flexibility. They use a radically different architecture than RDBMS, by exploiting (rather than embedding) a distributed file system such as Google File System (GFS) or Hadoop Distributed File System (HDFS), to store and manage data in a highly fault-tolerant manner. They tend to rely on a more specific data model, e.g. key-value store such Google Bigtable, Hadoop Hbase or Apache CouchDB) with a simple set of operators easy to use from a programming language. For instance, to address the requirements of social network applications, new solutions rely on a graph data model and graph-based operators. User-defined functions also allow for more specific data processing. MapReduce is a good example of generic parallel data processing framework, on top of a distributed file system (GFS) or HDFS). It supports a simple data model (sets of (key, value) pairs), which allows user-defined functions (map and reduce). Although quite successful among developers, it is relatively low-level and rigid, leading to custom user code that is hard to maintain and reuse. In Zenith, we exploit or extend these NOSQL technologies to fit our needs for scientific workflow management and scalable data analysis.

3.4. Uncertain Data Management

Data uncertainty is present in many scientific applications. For instance, in the monitoring of plant contamination by INRA teams, sensors generate periodically data which may be uncertain. Instead of ignoring (or correcting) uncertainty, which may generate major errors, we need to manage it rigorously and provide support for querying.

To deal with uncertainty, there are several approaches, e.g. probabilistic, possibilistic, fuzzy logic, etc. The *probabilistic approach* is often used by scientists to model the behavior of their underlying environments. However, in many scientific applications, data management and uncertain query processing are not integrated, i.e. the queries are usually answered using ad-hoc methods after doing manual or semi-automatic statistical treatment on the data which are retrieved from a database. In Zenith, we aim at integrating scientific data management and query processing within one system. This should allow scientists to issue their queries in a query language without thinking about the probabilistic treatment which should be done in background in order to answer the queries. There are two important issues which any PDBMS should address: 1) how to represent a probabilistic database, i.e. data model; 2) how to answer queries using the chosen representation, i.e. query evaluation.

One of the problems on which we focus is *scalable query processing* over uncertain data. A naive solution for evaluating probabilistic queries is to enumerate all possible worlds, i.e. all possible instances of the database, execute the query in each world, and return the possible answers together with their cumulative probabilities. However, this solution can not scale up due to the exponential number of possible worlds which a probabilistic database may have. Thus, the problem is quite challenging, particularly due to exponential number of possibilities that should be considered for evaluating queries. In addition, most of our underlying scientific applications are not centralized; the scientists share part of their data in a *P2P* manner. This distribution of data makes very complicated the processing of probabilistic queries. To develop efficient query processing techniques for distributed scientific applications, we can take advantage of two main distributed technologies: *P2P* and *Cloud*. Our research experience in P2P systems has proved us that we can propose scalable solutions for many data management problems. In addition, we can use the cloud parallel solutions, e.g. MapReduce, to parallelize the task of query processing, when possible, and answer queries of scientists in reasonable execution times. Another challenge for supporting scientific applications is uncertain data integration. In addition to managing the uncertain data for each user, we need to integrate uncertain data from different sources. This requires revisiting traditional data integration in major ways and dealing with the problems of uncertain mediated schema generation and uncertain schema mapping.

3.5. Metadata Integration

Nowdays, scientists can rely on web 2.0 tools to quickly share their data and/or knowledge (e.g. ontologies of the domain knowledge). Therefore, when performing a given study, a scientist would typically need to access and integrate data from many data sources (including public databases). To make high numbers of scientific data sources easily accessible to community members, it is necessary to identifying semantic correspondences between metadata structures or models of the related data sources. The main underlying task is called matching, which is the process of discovering semantic correspondences between metadata structures such as database schema and ontologies. Ontology is a formal and explicit description of a shared conceptualization in term of concepts (i.e., classes, properties and relations). For example, the matching may be used to align gene ontologies or anatomical metadata structures.

To understand a data source content, metadata (data that describe the data) is crucial. Metadata can be initially provided by the data publisher to describe the data structure (e.g. schema), data semantics based on ontologies (that provide a formal representation of the domain knowledge) and other useful information about data provenance (publisher, tools, methods, etc.). Scientific metadata is very heterogeneous, in particular because of the great autonomy of the underlying data sources, which leads to a large variety of models and formats. The high heterogeneity makes the matching problem very challenging. Furthermore, the number of ontologies and their size grow fastly, so does their diversity and heterogeneity. As a result, schema/ontology matching has become a prominent and challenging topic [4].

3.6. Data Mining

Data mining provides methods to discover new and useful patterns from very large sets of data. These patterns may take different forms, depending on the end-user's request, such as:

- **Frequent itemsets and association rules [1]**. In this case, the data is usually a table with a high number of rows and the algorithm extracts correlations between column values. This problem was first motivated by commercial and marketing purposes (*e.g.* discovering frequent correlations between items bought in a shop, which could help selling more). A typical example of frequent itemset from a sensor network in a smart building would say that "in 20% rooms, the door is closed, the room is empty, and lights are on."
- **Frequent sequential pattern extraction.** This problem is very similar to frequent itemset mining, but in this case, the order between events has to be considered. Let us consider the smart-building example again. A frequent sequence, in this case, could say that "in 40% rooms, lights are on at time i, the room is empty at time i+j and the door is closed at time i+j+k". Discovering frequent sequences has become a crucial need in marketing, but also in security (detecting network intrusions for instance) in usage analysis (web usage is one of the main applications) and any domain where data arrive in a specific order (usually given by timestamps).
- **Clustering [10]**. The goal of clustering algorithms is to group together data that have similar characteristics, while ensuring that dissimilar data will not be in the same cluster. In our example of smart buildings, we would find clusters of rooms, where offices will be in one category and copy machine rooms in another one because of their characteristics (hours of people presence, number of times lights are turned on and off, etc.).

One of the main problems for data mining methods recently was to deal with data streams. Actually, data mining methods have first been designed for very large data sets where complex algorithms of artificial intelligence were not able to complete within reasonable time responses because of data size. The problem was thus to find a good trade-off between time response and results relevance. The patterns described above well match this trade-off since they both provide interesting knowledge for data analysts and allow algorithm having good time complexity on the number of records. Itemset mining algorithms, for instance, depend more on the number of columns (for a sensor it would be the number of possible items such as temperature, presence, status of lights, etc.) than the number of lines (number of sensors in the network). However, with the ever growing size of data and their production rate, a new kind of data source has recently emerged as data streams. A data stream is a sequence of events arriving at high rate. By "high rate", we usually admit that traditional data mining methods reach their limits and cannot complete in real-time, given the data size. In order to extract knowledge from such streams, a new trade-off had to be found and the data mining community has investigated approximation methods that could allow maintaining a good quality of results for the above patterns extraction.

For scientific data, data mining now has to deal with new and challenging characteristics. First, scientific data is often associated to a level of uncertainty (typically, sensed values have to be associated to the probability that this value is correct or not). Second, scientific data might be extremely large and need cloud computing solutions for their storage and analysis. Eventually, we will have to deal with high dimension and heterogeneous data.

3.7. Content-based Information Retrieval

Today's technologies for searching information in scientific data mainly rely on relational DBMS or text based indexing methods. However, content-based information retrieval has progressed much in the last decade and is now considered as one of the most promising for future search engines. Rather than restricting search to the use of metadata, content-based methods attempt to index, search and browse digital objects by means of signatures describing their actual content. Such methods have been intensively studied in the multimedia community to allow searching the massive amount or raw multimedia documents created every day (e.g. 99% of web data are audio-visual content with very sparse metadata). Successful and scalable content-based methods have been proposed for searching objects in large image collections or detecting copies in huge video archives. Besides multimedia contents, content-based information retrieval methods recently started to be studied on more diverse data such as medical images, 3D models or even molecular data. Potential applications in scientific data management are numerous. First of all, to allow searching the huge collections of scientific images (earth observation, medical images, botanical images, biology images, etc.) but also to browse large datasets of experimental data (e.g. multisensor data, molecular data or instrumental data). Despite recent progress, scalability remains a major issue, involving complex algorithms (such as similarity search, clustering or supervised retrieval), in high dimensional spaces (up to millions of dimensions) with complex metrics (Lp, Kernels, sets intersections, edit distances, etc.). Most of these algorithms have linear, quadratic or even cubic complexities so that their use at large scale is not affordable without consistent breakthrough. In Zenith, we plan to investigate the following challenges:

- **High-dimensional similarity search**. Whereas many indexing methods were designed in the last 20 years to retrieve efficiently multidimensional data with relatively small dimensions, high-dimensional data have been more challenging due to the well-known dimensionality curse. Only recently have some methods appeared that allow approximate Nearest Neighbors queries in sub-linear time, in particular, Locality Sensitive Hashing methods which offer new theoretical insights in high-dimensional Euclidean spaces and proved the interest of random projections. But there are still some challenging issues that need to be solved including efficient similarity search in any kernel or metric spaces, efficient construction of knn-graphs or relational similarity queries.
- **Large-scale supervised retrieval**. Supervised retrieval aims at retrieving relevant objects in a dataset by providing some positive and/or negative training samples. To solve such task, there has been a focused interest on using Support Vector Machine's (SVM) that offer the possibility to construct generalized, non-linear predictors in high-dimensional spaces using small training sets. The prediction time complexity of these methods is usually linear in dataset size. Allowing hyperplane similarity queries in sub-linear time is for example a challenging research issue. A symmetric problem in supervised retrieval consists in retrieving the most relevant object categories that might contain a given query object, providing huge labeled datasets (up to millions of classes and billions of objects) and very few objects per category (from 1 to 100 objects). SVM methods that are formulated as quadratic programming with cubic training time complexity and quadratic space complexity are clearly not usable. Promising solutions to such problems include hybrid supervised-unsupervised methods and supervised hashing methods.

- **P2P content-based retrieval**. Content-based P2P retrieval methods appeared recently as a promising solution to manage masses of data distributed over large social networks, particularly when the data cannot be centralized for privacy or cost reasons (which is often the case in scientific social networks, e.g. botanist social networks). However, current methods are limited to very simple similarity search paradigms. In Zenith, we will consider more advanced P2P content-based retrieval and mining methods such as k-nn graphs construction, large-scale supervised retrieval or multi-source clustering.

4. Application Domains

4.1. Application Domains

The application domains covered by Zenith are very wide and diverse, as they concern data-intensive scientific applications, i.e. most scientific applications. Since the interaction with scientists is crucial to identify and tackle data management problems, we are dealing primarily with application domains for which Montpellier has an excellent track record, i.e. agronomy, environmental science, life science, with scientific partners like INRA, CIRAD and CEMAGREF. However, we are also addressing other scientific domains (e.g. astronomy, oil extraction) through our international collaborations (e.g. in Brazil).

Let us briefly illustrate three representative examples of scientific applications on which we have been working on.

- **Pesticide reduction**. In a pesticide reduction application, with CEMAGREF, we plan to work on sensor data for plant monitoring. Sensors are used to observe the development of diseases and insect attacks in the agricultural farms, aiming at using pesticides only when necessary. The sensors periodically send to a central system their data about different measures such as plants contamination, temperature or moisture level. A decision support system analyzes the sent data, and triggers a pesticide treatment only when needed. However, the data sent by sensors are not entirely certain. The main reasons for uncertainty are the effect of climate events on sensors, e.g. rain, unreliability of the data transmission media, fault in sensors, etc. This requires to deal with uncertain data in modeling and querying to be used for data analysis and data mining.
- **Botanical data sharing**. Botanical data is highly decentralized and heterogeneous. Each actor has its own expertise domain, hosts its own data, and describes them in a specific format. Furthermore, botanical data is complex. A single plant's observation might include many structured and unstructured tags, several images of different organs, some empirical measurements and a few other contextual data (time, location, author, etc.). A noticeable consequence is that simply identifying plant species is often a very difficult task; even for the botanists themselves (the so-called taxonomic gap). Botanical data sharing should thus speed up the integration of raw observation data, while providing users an easy and efficient access to integrated data. This requires to deal with social-based data integration and sharing, massive data analysis and scalable content-based information retrieval. We address this application in the context of the French initiative Pl@ntNet, with CIRAD and IRD.
- **Deepwater oil exploitation**. An important step in oil exploitation is pumping oil from ultra-deepwater from thousand meters up to the surface through long tubular structures, called risers. Maintaining and repairing risers under deep water is difficult, costly and critical for the environment. Thus, scientists must predict risers fatigue based on complex scientific models and observed data for the risers. Risers fatigue analysis requires a complex workflow of data-intensive activities which may take a very long time to compute. A typical workflow takes as input files containing riser information, such as finite element meshes, winds, waves and sea currents, and produces result analysis files to be further studied by the scientists. It can have thousands of input and output files and tens of activities (e.g. dynamic analysis of risers movements, tension analysis, etc.). Some activities, e.g. dynamic analysis, are repeated for many different input files, and depending on the

mesh refinements, each single execution may take hours to complete. To speed up risers fatigue analysis requires parallelizing workflow execution, which is hard to do with existing SWfMS. We address this application in collaboration with UFRJ, and Petrobras.

These three application examples illustrate the diversity of requirements and issues which we are addressing with our scientific application partners (CIRAD, INRA, CEMAGREF, ...). To further validate our solutions and extend the scope of our results, we also want to foster industrial collaborations, even in non scientific applications, provided that they exhibit similar challenges.

5. Software

5.1. WebSmatch (Web Schema Matching)

Participants: Zohra Bellahsène, Emmanuel Castanier, Rémi Coletta, Duy Hoa Ngo, Patrick Valduriez [contact].

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

In the context of the Action de Développement Technologique (ADT) started in october 2010, WebSmatch is a flexible, open environment for discovering and matching complex schemas from many heterogeneous data sources over the Web. It provides three basic functions: (1) metadata extraction from data sources; (2) schema matching (both 2-way and n-way schema matching), (3) schema clustering to group similar schemas together. WebSmatch is being delivered through Web services, to be used directly by data integrators or other tools, with RIA clients. Implemented in Java, delivered as Open Source Software (under LGPL) and protected by a deposit at APP (Agence de Protection des Programmes). WebSmatch is being used by Datapublica and CIRAD to integrate public data sources.

5.2. SON (Shared-data Overlay Network)

Participants: Ayoub Ait Lahcen, Fady Draidi, Esther Pacitti, Didier Parigot [contact], Patrick Valduriez, Guillaume Verger.

URL: http://www-sop.inria.fr/teams/zenith/SON

SON is an open source development platform for P2P networks using web services, JXTA and OSGi. SON combines three powerful paradigms: components, SOA and P2P. Components communicate by asynchronous message passing to provide weak coupling between system entities. To scale up and ease deployment, we rely on a decentralized organization based on a DHT for publishing and discovering services or data. In terms of communication, the infrastructure is based on JXTA virtual communication pipes, a technology that has been extensively used within the Grid community. Using SON, the development of a P2P application is done through the design and implementation of a set of components. Each component includes a technical code that provides the component services and a code component that provides the component logic (in Java). The complex aspects of asynchronous distributed programming (technical code) are separated from code components and automatically generated from an abstract description of services (provided or required) for each component by the component generator.

5.3. P2Prec (P2P recommendation service)

Participants: Fady Draidi, Esther Pacitti [contact], Didier Parigot, Guillaume Verger.

URL: http://p2prec.gforge.inria.fr

P2Prec is recommendation service for P2P content sharing systems that exploits users social data. To manage users social data, we rely on Friend-Of-A-Friend (FOAF) descriptions. P2Prec has a hybrid P2P architecture to work on top of any P2P content sharing system. It combines efficient DHT indexing to manage the users FOAF files with gossip robustness to disseminate the topics of expertise between friends. P2Prec is implemented in java using the Data-Shared Overlay Network (SON) infrastructure which is the basis for the ANR DataRing project.

5.4. ProbDB (Probabilistic Database)

Participants: Reza Akbarinia [contact], Patrick Valduriez, Guillaume Verger.

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

ProbDB is a probabilistic data management system to manage uncertain data on top of relational DBMSs. One of the main features of the prototype is its portability; that means with a minimum effort it can be implemented over any DBMS. In ProbDB, we take advantage of the functionalities provided by almost all DBMSs, particularly the query processing functions. It is implemented in Java on top of PostgreSQL.

5.5. SnoopIm

Participants: Julien Champ [contact], Alexis Joly.

SnoopIm is a content-based search engine allowing to retrieve small visual patterns or objects in large collections of pictures (such as logos on clothes, road signs in the background, paintings on walls, etc.) and to derive statistics from them (frequency, visual cover, size variations, etc.). Query objects to be searched can be either selected from the collection of photos or from an external picture (by simply providing its URL). The web application allows online search of multiple users and has a cache feature to speed-up the processing of seen queries. It is implemented in Javascript on top of a C++ library developed in collaboration with INA'sup (http://www.ina-sup.com/).

5.6. SimJoin (Distributed Approximate Similarity Join)

Participant: Alexis Joly [contact].

SimJoin is a distributed software for the efficient computation of the full approximate k-nn graph of large collections of high-dimensional features. It is developed within a MapReduce framework and is therefore easily portable to large cloud computing plateform. It is based on recent theoretic contributions related to locality preserving hash functions [34]. Its first main feature is to allow splitting a large collection of high-dimensional features into highly balanced pages that preserve locality according to any given similarity kernel. Its second main feature is to build in $O(n^{1+\gamma})$ operations a candidate set of item pairs that approximate the theoretic knn-graph with high recall. This software is developed in collaboration with INRIA Imedia.

6. New Results

6.1. Data and Metadata Management

6.1.1. Uncertain Data Management

Participants: Reza Akbarinia, Patrick Valduriez, Guillaume Verger.

Data uncertainty in scientific applications can be due to many different reasons: incomplete knowledge of the underlying system, inexact model parameters, inaccurate representation of initial boundary conditions, inaccuracy in equipments, etc. For instance, in the monitoring of plant contamination, sensors generate periodically data which may be uncertain. Instead of ignoring (or correcting) uncertainty, which may generate major errors, we need to manage it rigorously and provide support for querying.

In [46], we address the problem of aggregate queries that return possible sum values and their probabilities. This kind of query which, we call ALL-SUM, is also known as sum probability distribution. The results of ALL-SUM can be used for answering many other type of queries over probabilistic data. In general, the problem of ALL-SUM query execution is NP-complete. We propose pseudo-polynomial algorithms that are efficient in many practical applications, e.g. when the aggr attribute values are small integers or real numbers with small precision, i.e. small number of digits after decimal point. These cases cover many practical attributes, e.g. temperature, blood pressure, needed human recourses per patient in medical applications.

We have started to develop a probabilistic database prototype, called ProbDB (Probabilistic Database), on top of an RDBMS. ProbDB divides a query into two parts: probabilistic and deterministic (i.e. non probabilistic). The deterministic part is executed by the underlying RDBMS, and the rest of work is done by our probabilistic query processing algorithms that are executed over the data returned by the RDBMS. In [51], we demonstrated the efficient execution of aggregate queries with the first version of ProbDB.

6.1.2. Metadata Integration

Participants: Zohra Bellahsène, Rémi Coletta, Duy Hoa Ngo.

Due to the various types of heterogeneity of ontologies, ontology matching must exploit many features of ontology elements in order to improve matching quality. For this purpose, numerous similarity metrics have been proposed to deal with ontology semantics at different levels: elements level, structural level and instance level.

Elements level metrics can be categorized in three groups: (1) terminological, (2) structural and (3) semantic. Metrics of the first group exploit text features such as names, labels and comments to compute the similarity score between entities. Whereas metrics of the last two groups exploit the hierarchy and semantic relationship features. Our approach consists in first using terminological metrics. Then, during the matching process, mappings discovered by terminological metrics are used as input mappings to other metrics of the second and third groups. Obviously, the more precise results terminological metrics are, the more accurate results structural and semantic metrics have.

However, finding a good combination of different metrics is very difficult and time consuming. We proposed YAM++ (not Yet Another Matcher), an approach that uses machine learning to combine similarity metrics. Our main contributions are: the definition of new metrics dealing with terminological and context profile features of entities in ontologies [37], and the use of a decision tree model to combine similarity metrics [38].

To improve matching quality of YAM++, we exploit instances accompanying ontologies. We then apply similarity flooding propagation algorithm to discover more semantic mappings. At the 2011 competition of the Ontology Alignment Evaluation Initiative (http://oaei.ontologymatching.org), YAM++ achieved excellent results: first position on the Conference track and second position on the Benchmark track [39].

6.2. Data and Process Sharing

6.2.1. Social-based P2P Data Sharing

Participants: Hinde Bouziane, Michèle Cart, Esther Pacitti, Didier Parigot, Guillaume Verger.

This work focuses on P2P content recommendation for on-line communities. In [20], we propose P2Prec, a recommendation service for P2P content sharing systems that exploits users' social data. Given a query, P2PRec finds peers that can recommend high quality documents that are relevant for the query. A document is relevant to a query if it covers the same topics. It is of high quality if relevant peers have rated it highly. P2PRec finds relevant peers through a variety of mechanisms including advanced content-based and collaborative filtering. The topics each peer is interested in are automatically calculated by analyzing the documents the peer holds. Peers become relevant for a topic if they hold a certain number of highly rated documents on this topic. To efficiently disseminate information about peers' topics and relevant peers, we proposed new semantic-based gossip protocols. In our experimental evaluation, using the TREC09 dataset, we showed that using semantic gossip increases recall by a factor of 1.6 compared to well-known random gossiping. Furthermore, P2Prec has the ability to get reasonable recall with acceptable query processing load and network traffic. P2Prec was demonstred in [31] and [47].

In [30], we exploit social relationships between users as a parameter to increase the trust of recommendation. We propose a novel P2P recommendation approach (called F2Frec) that leverages content and social-based recommendation by maintaining a P2P and friend-to-friend network. This network is used as a basis to provide useful and high quality recommendations. Based on F2Frec, we propose new metrics, such as usefulness and similarity (among users and their respective friend network). We define our proposed metrics based on users' topic of interest and relevant topics that are automatically extracted from the contents stored by each user. Our experimental evaluation, using the TREC09 dataset and Wiki vote social network, shows the benefits of our approach compared to anonymous recommendation. In addition, we show that F2Frec increases recall by a factor of 8.8 compared with centralized collaborative filtering.

6.2.2. Satisfaction-based Query Replication

Participant: Patrick Valduriez.

In a large-scale Internet-based distributed, participants (consumers and providers) who are willing to share data are typically autonomous, i.e. they may have special interests towards queries and other participants' data. In this context, a way to avoid a participant to voluntarily leave the system is satisfying its interests when allocating queries. However, participants' satisfaction may also be negatively affected by the failures of other participants. Query replication can deal with providers failures, but, it is challenging because of autonomy: it cannot only quickly overload the system, but also dissatisfy participants with uninteresting queries. Thus, a natural question arises: should queries be replicated? If so, which ones? and how many times?

In [25], we answer these questions by revisiting query replication from a satisfaction and probabilistic point of view. We propose a new algorithm, called S b QR, that decides on-the-fly whether a query should be replicated and at which rate. As replicating a large number of queries might overload the system, we propose a variant of our algorithm, called S b QR+. The idea is to voluntarily fail to allocate as many replicas as required by consumers for low critical queries so as to keep resources for high critical queries during query-intensive periods. Our experimental results demonstrate that our algorithms significantly outperform the baseline algorithms from both the performance and satisfaction points of view. We also show that our algorithms automatically adapt to the criticality of queries and different rates of participant failures.

6.2.3. View Selection in Scientific Data Warehousing

Participants: Zohra Bellahsène, Rémi Coletta, Imen Mami.

Scientific data generate large amounts of data which have to be collected and stored for analytical purpose. One way to help managing and analyzing large amounts of data is data warehousing, whereby views over data are materialized. However, view selection is an NP-hard problem because of many parameters: query cost, view maintenance cost and storage space. In [36], we propose a new solution based on constraint programming, which has proven efficient at solving combinatorial problems. This allows using a constraint programming solver to set up the search space by identifying a set of views that minimizes the total query cost. We address view selection under two cases: (1) only the total view maintenance cost needs be minimed, assuming unlimited storage space (meaning that it is not a critical resource anymore); (2) both storage space and maintenance cost must be mimized. We implemented our approach and compared it with a randomized method (i.e., genetic algorithm). We experimentally show that our approach provides better performance resulting from evaluating the quality of the solutions in terms of cost savings. Furthermore, our approach scales well with the query workload.

6.2.4. Scientific Workflow Management

Participants: Ayoub Ait Lahcen, Eduardo Ogasawara, Didier Parigot, Patrick Valduriez.

Scientific workflows have emerged as a basic abstraction for structuring and executing scientific experiments in computational environments. In many situations, these workflows are computationally and data-intensive, thus requiring execution in large-scale parallel computers. However, parallelization of scientific workflows remains low-level, ad-hoc and laborintensive, which makes it hard to exploit optimization opportunities.

To address this problem, we propose in [23] an algebraic approach (inspired by relational algebra) and a parallel execution model that enable automatic optimization of scientific workflows. With our scientific workflow algebra, data is uniformly represented by relations and workflow activities are mapped to operators that have data aware semantics. Our workflow execution model is based on the concept of activity activation, which enables transparent distribution and parallelization of activities;

We conducted a thorough validation of our approach using both a real oil exploitation application and synthetic data scenarios. The experiments were run in Chiron, a data-centric scientific workflow engine implemented at UFRJ to support our algebraic approach. Our experiments demonstrate performance improvements of up to 226% compared to an ad-hoc workflow implementation. This work was done in the context of the Equipe Associée Sarava and the CNPq-INRIA project DatLuge.

In the context of SON, we also proposed a declarative workflow language based on service/activity rules [41]. This language makes it possible to infer a dependency graph for SON applications that provides for automatic parallelization.

6.3. Scalable Data Analysis

6.3.1. Massive Graph Management

Participant: Patrick Valduriez.

Traversing massive graphs as efficiently as possible is essential for many scientific applications. Many common operations on graphs, such as calculating the distance between two nodes, are based on the Breadth First Search (BFS) traversal. However, because of the exhaustive exploration of all the nodes and edges of the graph, this operation might be very time consuming. A possible solution is partitioning the graph among the nodes of a shared-nothing parallel system. However, partitioning a graph and keeping the infor- mation regarding the location of vertices might be unrealistic for massive graphs because of much inter-node communication. In [28], we propose ParallelGDB, a new graph database system based on specializing the local caches of any node in this system, providing a better cache hit ra- tio. ParallelGDB uses a random graph partitioning, avoid- ing complex partition methods based on the graph topology, that usually require managing extra data structures. This proposed system provides an efficient environment for dis- tributed graph databases.

6.3.2. Top-k Query Processing in Unstructured P2P Systems

Participants: Reza Akbarinia, William Kokou Dedzoe, Patrick Valduriez.

Top-k query processing techniques are useful in unstructured P2P systems to avoid overwhelming users with too many results and provide them with the best ones. However, existing approaches suffer from long waiting times, because top-k results are returned only when all queried peers have finished processing the query. As a result, response time is dominated by the slowest queried peer. We proposed to revisit the problem of top-k query processing.

In [29] we address the problem of reducing user waiting time of top-k query processing, in the case of unstructured P2P systems with overloaded peers. We propose a new algorithm, called QUAT, in which each peer maintains a semantic description of its local data and the semantic descriptions of its neighborhood (i.e. the semantic descriptions of data owned locally by its direct neighbors and data owned locally by these neighbors direct neighbors). These semantic descriptions allow peers to prioritize the queries that can provide high quality results, and to forward them in priority to the neighbors that can provide high quality answers. We validated our solution through a thorough experimental evaluation using a real-world dataset. The results show that QUAT significantly outperforms baseline algorithms by returning faster the final top-k results to users.

6.3.3. Top-k Query Processing Over Sorted Lists

Participants: Reza Akbarinia, Esther Pacitti, Patrick Valduriez.

The problem of answering top-k queries can be modeled as follows. Suppose we have m lists of n data items such that each data item has a local score in each list and the lists are sorted according to the local scores of their data items. Each data item has an overall score computed based on its local scores in all lists using a given scoring function. Then, the problem is to find the k data items whose overall scores are the highest. This problem model is a general model for top-k queries in many centralized, distributed and P2P applications. For example, in IR systems one of the main problems is to find the top-k documents whose aggregate rank is the highest wrt. some given keywords. To answer this query, the solution is to have for each keyword a ranked list of documents, and return the k documents whose aggregate rank in all lists are the highest.

In [16], we propose an extension of our best position algorithms (BPA) which had been proposed for top-k query processing over sorted lists model.The BPA algorithms have been shown to be more efficient than the well known TA Algorithm. We propose several techniques using different data structures for managing best positions that are crucial for efficient execution of top-k algorithms. We also provide a complete discussion on the instance optimality of TA algorithm (TA was considered so far as optimal over any database of sorted lists). We illustrate that, the existence of deterministic algorithms such as BPA shows that if we are aware of positions of seen data, then one of the main arguments used for proving the instance optimality of TA is invalidated. Therefore, in this case the proof of TA's instance optimality is incorrect, and must be revisited.

6.3.4. Satellite Image Mining

Participant: Florent Masseglia.

Satellite Image Time Series (SITS) provide us with precious information on land cover evolution. By studying SITS, we can both understand the changes of specific areas and discover global phenomena that spread over larger areas. Changes that can occur throughout the sensing time can spread over very long periods and may have different start time and end time depending on the location, which complicates the mining and the analysis of series of images. In [45], we propose a frequent sequential pattern mining method for SITS analysis. Designing such a method called for important improvements on the data mining principles. First, the search space in SITS is multi-dimensional (the radiometric levels of different wavelengths correspond to infra-red, red, etc.). Furthermore, the non evolving regions, which are the vast majority and overwhelm the evolving ones, challenge the discovery of these patterns. Our framework enables discovery of these patterns despite these constraints and characteristics. We introduce new filters in the mining process to yield important reductions in the search space by avoiding consecutive occurences of similar values in the sequences. Then, we propose visualization techniques for results analysis (where modified regions are highlighted). Experiments carried out on a particular dataset showed that our method allows extracting repeated, shifted and distorted temporal behaviors. The flexibility of this method makes it possible to capture complex behaviors from multi-source, noisy and irregularly sensed data.

6.3.5. Distributed Approximate Similarity Join

Participant: Alexis Joly.

Efficiently constructing the KNN-graph of large and high dimensional feature datasets is crucial for many data intensive applications involving feature-rich objects, such as image features, text features or sensor's features. In this work we investigate the use of high dimensional hashing methods for efficiently approximating the full knn graph of large collections, in particular, in distributed environments. We first analyzed and experimented what seems to be the most intuitive hashing-based approach: constructing several Locality Sensitive Hashing (LSH) tables in parallel and computing the frequency of all emitted collisions. We show that balancing issues of classical LSH functions strongly affect the performance of this approach. On the other side, we show that using an alternative data-dependent hashing function (RMMH), that we introduced recently [34], can definitely change that conclusion. The main originality of RMMH hash function family is that it is based on randomly trained classifiers, allowing to learn random and balanced splits of the data instead of using random splits of the feature space as in LSH. We show that the hash tables constructed through RMMH are much more balanced and that the number of emitted collisions can be strongly reduced without degrading quality. In the end, our hashing-based filtering algorithm of the all-pairs graph is two orders of magnitude faster than the one based on LSH. An efficient distributed implementation of the method was implemented within the MapReduce framework (and is the basis of the SimJoin prototype). This work is done in the context of the supervision of a PhD student working at INRIA Imedia (Riadh Mohamed Trad).

6.3.6. Visual objects mining

Participant: Alexis Joly.

State-of-the-art content-based object retrieval systems have demonstrated impressive performance in very large image datasets. These methods, based on fine local descriptions and efficient matching techniques, can detect accurately very small rigid objects with unambiguous semantics such as logos, buildings, manufactured objects, posters, etc. Mining such small objects in large collections is however difficult. Constructing a full local matching graph with a naïve approach would indeed require to probe all candidate query leading to an intractable algorithm complexity. In this work, we first introduce an adaptive weighted sampling scheme, starting with some prior distribution and iteratively converging to unvisited regions [35]. We show that the proposed method allows to discover highly interpretable visual words while providing excellent recall and image representativity. We then focused on mining visual objects on top of the discovered visual words. We therefore developed an original shared nearest-neighbors clustering method, working directly on the generated bi-partite graph. This work is in the context of the supervision of two PhD students, one working jointly with INA and INRIA and who will join the Zenith team next year (Pierre Letessier), one working at INRIA Imedia (Amel Hamzaoui).

6.3.7. Visual-based plant species identification from crowdsourced data

Participant: Alexis Joly.

Inspired by citizen sciences, the main goal of this work is to speed up the collection and integration of raw botanical observation data, while providing to potential users an easy and efficient access to this botanical knowledge. We therefore designed and developed an original crowdsourcing web application dedicated to the access of botanical knowledge through automated identification of plant species by visual content. Technically, the first side of the application deals with content-based identification of plant leaves. Whereas state-of-the-art methods addressing this objective are mostly based on leaf segmentation and boundary shape features, we developed a new approach based on local features and large-scale matching. This approach obtained the best results within ImageCLEF 2011 plant identification benchmark [48]. The second side of the application deals with interactive tagging and allows any user to validate or correct the automatic determinations returned by the system. Overall, this collaborative system allows to enrich automatically and continuously the visual botanical knowledge and thus to increase progressively the accuracy of the automated identification. A demo of the developed application was presented at the ACM Multimedia conference [33]. This work is done in collaboration with INRIA Imedia and with the botanists of the AMAP UMR team (CIRAD). It is also closely related to a citizen science project around plant's identification that we develop with the support of the TelaBotanica social network.

7. Contracts and Grants with Industry

7.1. ANR OTMedia (2011-2013), 150 Keuros

Participants: Alexis Joly, Julien Champ.

The project OTMedia is with INA, AFP, Syllabs, LIA, and ELICO. The main objective is to create an integrated system allowing sociologists and decision makers to analysis information flows across different media sources including web sites, blogs, newspapers, radio and TV. Zenith addresses more specifically the research challenges related to the trans-media tracking of visual contents (images and videos).

7.2. CIFRE INA/INRIA (2011-2013)

Participants: Alexis Joly, Pierre Letessier.

This CIFRE contract with INA allows funding a 3-years PhD (Pierre Letessier). This PhD addresses research challenges related to content-based mining of visual objects in large collections.

7.3. PREDIT EPILOG (2009-2011, 60Keuros)

Participant: Patrick Valduriez.

The project EPILOG (Etude des technologies Pair-à-pair pour la collaboration Interentreprises dans la chaîne LOGistique) involves Euxenis SAS and RISC Solutions d'Assurances. The objective is to provide support for collaboration and supply chain management among partner enterprises in the retail industry The approach we validated in the project is P2P. We also addressed the research issues associated with the definition of the P2P network for supply chain management, with autonomous partners with various interests, the modeling of information exchanged during transactions and query processing in the P2P network.

7.4. Data Publica (2010-2011)

Participants: Emmanuel Castanier, Rémi Coletta, Patrick Valduriez.

Data Publica (http://www.data-publica.com) is a startup providning a web portal for open data which can be public, private, free or charged. We collaborate with Data Publica though our WebSmatch technology on technologies for automatic schema extraction and matching from high numbers of data sources. A first contribution has been the development of an Excel extraction component based on machine learning techniques.

8. Partnerships and Cooperations

8.1. Regional Initiatives

8.1.1. Labex NUMEV, Montpellier

In the context of the Excellence Initiative of the MENRT, we are participating in the Laboratory of Excellence (labex) NUMEV (Digital and Hardware Solutions, Modelling for the Environment and Life Sciences) headed by University of Montpellier 2 in partnership with CNRS, University of Montpellier 1, and INRIA. NUMEV seeks to harmonize the approaches of hard sciences and life and environmental sciences in order to pave the way for an emerging interdisciplinary group with an international profile. The NUMEV project is decomposed in four complementary research themes: Modelling, Algorithms and computation, Scientific data (processing, integration, security), Model-Systems and measurements. Patrick Valduriez heads the theme on scientific data.

8.2. National Actions

8.2.1. ANR VERSO DataRing(2008-2012, 200Keuros)

Participants: Reza Akbarinia, Fady Draidi, Mohamed Jawad, Esther Pacitti, Guillaume Verger, Patrick Valduriez [contact].

URL: http://www-sop.inria.fr/teams/zenith/dataring

The DataRing project, headed by P. Valduriez, involves the Leo project-team (INRIA Saclay Ile de France), LIG, LIRMM and Telecom ParisTech. The objective is to address the problem of data sharing for online communities, such as social networks (e.g. sites like MySpace and Facebook) and professional communities (e.g. research communities, online technical support groups) which are becoming a major killer application of the web. The project addresses this problem by organizing community members in a peer-to-peer (P2P) network ring across distributed data source owners where each member can share data with the others through a P2P overlay network. In this project, we study the following problems: schema matching, query processing with data uncertainty, data indexing and caching, data privacy and trust. To validate our approach, we develop services based on our prototypes WebSmatch, SON, P2Prec and ProbDB.

8.2.2. ANR STAMP (2008-2011, 60Keuros)

Participants: Ayoub Ait Lahcen, Didier Parigot [contact], Guillaume Verger.

URL: http://www-sop.inria.fr/members/Didier.Parigot/pmwiki/Didier/index.php/Contracts/STAMP

The STAMP project is with CIRAD, INRA, CNRS and U. Paris-Est. The overall objective is to overcome present limitations of dynamic landscape modeling with spatial, temporal and multi-scale primitives. Our approach is to explore new spatial and temporal primitives and the potential benefits that recent advances in meta-modeling and Model-Driven Engineering can bring into the field of landscape studies. We contribute, together with CIRAD and U. Paris-Est, to the definition of a domain-specific language, called Ocelet.

8.3. International Initiatives

8.3.1. INRIA Associate Teams

8.3.1.1. SARAVA

Title: Sarava: P2P data sharing for online communities
INRIA principal investigator: Patrick Valduriez
International Partner:

Institution: Federal University of Rio de Janeiro (Brazil)
Laboratory: Instituto Alberto Luiz Coimbra

Duration: 2009 - 2011
See also: http://www-sop.inria.fr/teams/zenith/pmwiki/pmwiki.php/Sarava/Sarava
The general problem we address in Sarava is P2P data sharing for online communities, by offering a high-level network ring across distributed data source owners. The major advantage of P2P is a completely decentralized approach to data sharing which does not require centralized administration. Users may be in high numbers and interested in different kinds of collaboration and sharing their knowledge, ideas, experiences, etc. Data sources can be in high numbers, fairly autonomous, i.e. locally owned and controlled, and highly heterogeneous with different semantics and structures. What we need then is new, decentralized data management techniques that scale up while addressing the autonomy, dynamic behavior and heterogeneity of both users and data sources. In this context, we focus on two major problems: query processing with uncertain data and management of scientific workflows.

8.3.2. INRIA International Partners

We have regular scientific relationships with research laboratories in

- North America: Univ. of Waterloo (Tamer Özsu), McGill University (Bettina Kemme), Univ. of California, Santa Barbara (Divy Agrawal, Amr El Abbadi).
- Europe: Univ. of Madrid (Ricardo Jimenez-Periz).

8.3.3. Visits of International Scientists

Prof. Amr El Abbadi (UCSB) gave a seminar at LIRMM on cloud data management.

8.3.4. Participation In International Programs

We are involved in the following international actions:

- the CNPq-INRIA project DatLuge (Data & Task Management in Large Scale) with UFRJ (Marta Mattoso, Vanessa Braganholo, Alexandre Lima), LNCC, Rio de Janeiro (Fabio Porto), and UFPR, Curitiba (Eduardo Almeida) to work on large scale scientifc workflows;
- the PICASSO project Scaling GraphDB, with UPC, Barcelona ((Josep Lluis Larriba Pey and Victor Muntes Mulero) to work on very large graph database support.
- FAPERJ-INRIA project SwfP2Pcloud (Data-centric workflow management in hybrid P2P clouds, 2011-2013) with UFRJ (Marta Mattoso, Vanessa Braganholo, Alexandre Lima) and LNCC, Rio de Janeiro (Fabio Porto) to work on large scale scientific workflows in hybrid P2P clouds.
- EGIDE Osmoze project SECC (SErvices for Curricula Comparison, 2011-2012), with Riga Technical University (Janis Grundspenkis, Marité Kirikova) to work on automatic analysis and mapping of conceptual trees and maps acquired from digital documents.

9. Dissemination

9.1. Animation of the scientific community

Participation in the editorial board of scientific journals:

- VLDB Journal: P. Valduriez.
- Proceedings of the VLDB Endowment (PVLDB): E. Pacitti, P. Valduriez.
- Distributed and Parallel Databases, Kluwer Academic Publishers: P. Valduriez.
- Internet and Databases: Web Information Systems, Kluwer Academic Publishers: P. Valduriez.
- Journal of Information and Data Management, Brazilian Computer Society Special Interest Group on Databases: 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. 2011: Z. Bellahsène , P. Valduriez
- Int. Conf. on VLDB 2011: P. Valduriez
- Int. Conf. on Extending DataBase Technologies (EDBT), 2011: E. Pacitti, 2012: Z. Bellahsène, E. Pacitti
- IEEE Int. Conf. on Data Engineering (ICDE) 2011: P. Valduriez, 2012: Z. Bellahsène
- International Workshop on MapReduce and its Applications (MAPREDUCE), 2011: P. Valduriez.
- Journées Bases de Données Avancées (BDA), 2011: F. Masseglia, E. Pacitti, P. Valduriez
- Int. Conf. on Current Trends in Theory and Practice of Computer Science (SoftSem) 2011: P. Valduriez
- ACM Int. Conf. on Information and Knowledge Management (CIKM) 2011: Z. Bellahsène
- Int. Conf. on Cooperative Information systems (CoopIS) 2011: Z. Bellahsène
- Int. Conf. on Advanced Information Systems Engineering (CAiSE) 2011: Z. Bellahsène
- European Semantic Web Conference (ESWC) 2012: Z. Bellahsène
- ACM Symposium On Applied Computing (ACM SAC, data stream track), 2012: F. Masseglia
- Int. Conf. on Advances in Databases, Knowledge, and Data Applications (DBKDA), 2012: F. Masseglia
- Conférence Internationale Francophone sur l’Extraction et la Gestion de Connaissance (EGC), 2012: F. Masseglia
- IEEE Int. Conf. on Data Mining (ICDM), 2012: F. Masseglia
- Int. Conf. on Artificial Intelligence in Medicine (LEMEDS@AIME), 2011: F. Masseglia
- Int. Workshop on Multimedia Data Mining (MDM@KDD), 2011: F. Masseglia
- Int. Pacific-Asia Conf. on Knowledge Discovery and Data Mining (PAKDD), 2012: F. Masseglia
- IEEE Int. Conf. on Advanced Information Networking and Applications (AINA) 2012: H. Bouziane
- Rencontres Francophones sur les Aspects Algorithmiques de Télécommunications (AlgoTel), 2012: H. Bouziane
- NEM Summit 2011: A. Joly

The members of Zenith have always been strongly involved in organizing the French database research community, in the context of the I3 GDR and the BDA conference .

Patrick Valduriez was the panel chair of the 2011 IEEE Int. Conf. on Data Engineering (ICDE). He also organized the INRIA evaluation seminar of the theme Knowledge and Data Representation and Management.

Alexis Joly was a member of the organizing committee of the ImageCLEF evaluation forum (http://www.imageclef.org/2011). He organized a new evaluation task on plant images retrieval within the 2011 evaluation campaign of ImageCLEF. He was chair of a community networking session on the access to scientific multimedia data within CLEF 2011. He is also a member of the steering board of the CHORUS+ EU coordination action related to multimedia search engines (avmediasearch.eu) where he is leading a work package on the evaluation of multimedia retrieval technologies.

Esther Pacitti gave a tutorial on P2P techniques for decentralized applications at BDA 2011.

Patrick Valduriez was keynote speaker at DEXA 2011.

9.2. Teaching

Most permanent members of Zenith teach at the Licence and Master degree levels at UM2.

Zohra Bellahsène:

- IUT 1: Relational Databases , 100h, level L1, IUT, UM2
- IUT 2: Object-relational Databases, 60h, Level L2, IUT, UM2
- Master Computer Science : Data Integration, 15h, level M2, Faculty of Science, UM2

Hinde Bouziane:

Licence: Networks, 36h, level L3, Faculty of Science, UM2

Master MOCA: Networks and communication, 90h, level M1, Faculty of Science, UM2

Master IPS: Introduction to operating systems and networks, 36h, level M1, Faculty of Science, UM2

Master Research: Data sharing in P2P, 4,5h, level M2, Faculty of Science, UM2

Michelle Cart:

IG3 : Operating systems, 92h, level L3, Polytech'Montpellier, UM2

IG4 : Operating systems, introduction to distributed systems, 42h, level M1, Polytech'Montpellier, UM2

IG4, IG5 : student projects, 19h, Polytech'Montpellier, UM2

Rémi Coletta

IUT1: Relational Databases: 100h, level L1, IUT, UM2

Licence: Web Programming : 60h, level L3, IUT, UM2

Master Computer Science: Artificial Intelligence: 8h, level M1, Faculty of Science, UM2

Licence: Object-relational mapping with Hibernate, 10h, level L3, IUT, UM2

Master Computer Science: Constraint Programming, 20h, level M2, Faculty of Science, UM2

Esther Pacitti:

IG3: Database design, physical organization, 54h, level L3, Polytech'Montpellier, UM2

IG4: Networks, 42h, level M1, Polytech' Montpellier, UM2

IG4: Object-relational databases, 32h, level M1, Polytech' Montpellier, UM2

IG5: Distributed systems, virtualization, 27h, level M2, Polytech' Montpellier, UM2

Master Research: Data sharing in P2P, 4,5h, level M2, Faculty of Science, UM2

Patrick Valduriez:

Master Research: Data sharing in P2P, 12h, level M2, Faculty of Science, UM2

Professional: XML, 40h, level M2, Orsys Formation

PhD theses:

PhD : William Kokou Dedzoe, Top-k query processing in P2P virtual communities, 30 nov. 2011, Univ. Nantes, Advisors: Philippe Lamarre and Patrick Valduriez

PhD : Mohamed Jawad, Data privacy in P2P systems, 29 june 2011, Univ. Nantes, Advisor: Patrick Valduriez, Co-advisor: Patricia Serrano Alvarado

PhD : Mounir Tlili, A P2P infrastructure for data replication and reconciliation, 30 june 2011, Univ. Nantes, Advisor: Esther Pacitti, Co-advisor: Reza Akbarinia

PhD : Eduardo Ogasawara, An algebraic approach for data-centric scientific workflows, 19 dec. 2011, Universidade Federal de Rio de Janerio, Brazil, Advisors: Marta Mattoso and Patrick Valduriez

PhD in progress : Ayoub Ait Lahcen, Une Architecture Orientée Service (SOA) dynamique pour l'internet du futur, started Janv. 2009, Univ. Nice and Univ. Rabat Maroc, Advisor: Didier Parigot

PhD in progress : Fady Draidy, P2Prec: a social-based P2P recommendation system, started oct. 2008, UM2, Advisors: Esther Pacitti and Patrick Valduriez

PhD in progress : Vincenzo Gulisano, StreamCloud: a large-scale data stream management system, started oct. 2009, Universidad Politecnica de Madrid, Spain, Advisors: Ricardo Jimenez-Peris and Patrick Valduriez

PhD in progress : Pierre Letessier, Frequent Visual Objects Discovery in Multimedia Collections, started nov. 2009, Telecom ParisTech, Advisor: Nozha Boujemaa, co-advisors: Olivier Buisson and Alexis Joly

PhD in progress : Miguel Liroz, Massive Data Management for Scientific Applications, started oct. 2010, UM2, Advisors: Esther Pacitti and Patrick Valduriez, co-advisor: Reza Akbarinia

PhD in progress : Imen Mami, View selection: a constraint-programming based approach, started oct. 2009, UM2, Advisor: Zohra Bellahsène, co-advisor: Rémi Coletta

PhD in progress : Duy Hoa Ngo, A generic approach to ontology matching, started started oct. 2009, UM2, Advisor: Zohra Bellahsène, co-advisor: Rémi Coletta

PhD in progress : Toufik Sarni, Real-time support for software transactional memory in multicore systems, started oct. 2008, Univ. Nantes, Advisor: Patrick Valduriez, co-advisor: Audrey Queudet

PhD in progress : Maximilien Servajean, Decentralized and Personalized Recommendation Protocols for Content Sharing: application to phenotyping, started oct. 2011, UM2, Advisor: Esther Pacitti, co-advisors: Michèle Cart, Pascal Neveu

10. Bibliography

Major publications by the team in recent years

[1] P. PONCELET, F. MASSEGLIA, M. TEISSEIRE (editors). *Data Mining Patterns: New Methods and Applications*, Premier Reference Source, Idea Group, 2007, ISBN 978-1599041629, http://hal.inria.fr/lirmm-00365419/en.

[2] R. AKBARINIA, E. PACITTI, P. VALDURIEZ. *Best Position Algorithms for Efficient Top-k Query Processing*, in "Information Systems", 2011, vol. 36, n^{o} 6, p. 973-989, http://hal.inria.fr/lirmm-00607882/en.

[3] R. AKBARINIA, P. VALDURIEZ, G. VERGER. *Efficient Evaluation of SUM Queries Over Probabilistic Data*, in "IEEE Transactions on Knowledge and Data Engineering", 2012, To appear, http://hal.inria.fr/lirmm-00652293/en.

[4] Z. BELLAHSENE, A. BONIFATI, E. RAHM. *Schema Matching and Mapping*, Springer, March 2011, http://hal.inria.fr/lirmm-00581346/en.

[5] A. BENOIT, H. L. BOUZIANE, Y. ROBERT. *Optimizing the reliability of streaming applications under throughput constraints*, in "International Journal of Parallel Programming", 2011, vol. 39, n^{o} 5, p. 584-614 [*DOI :* 10.1007/s10766-011-0165-6], http://hal.inria.fr/inria-00574555/en.

[6] M. EL DICK, E. PACITTI, R. AKBARINIA, B. KEMME. *Building a Peer-to-Peer Content Distribution Network with High Performance, Scalability and Robustness*, in "Information Systems", 2011, vol. 36, n^{o} 2, p. 222-247, http://hal.inria.fr/lirmm-00607898/en.

[7] E. OGASAWARA, D. DE OLIVEIRA, P. VALDURIEZ, D. DIAS, F. PORTO, M. MATTOSO. *An Algebraic Approach for Data-Centric Scientific Workflows*, in "Proceedings of VLDB", 2011, vol. 4, n^{o} 11, p. 1328-1339, http://hal.inria.fr/hal-00640431/en.

[8] E. PACITTI, P. VALDURIEZ, M. MATTOSO. *Grid Data Management: Open Problems and New Issues*, in "Journal of Grid Computing", 2007, vol. 5, n^{o} 3, p. 273-281, http://hal.inria.fr/inria-00473481/en.

[9] J.-A. QUIANÉ-RUIZ, P. LAMARRE, P. VALDURIEZ. *A Self-Adaptable Query Allocation Framework for Distributed Information Systems*, in "The VLDB Journal", 2009, vol. 18, n^{o} 3, p. 649-674, http://hal.archives-ouvertes.fr/hal-00374999/fr/.

[10] C. ZHANG, F. MASSEGLIA, Y. LECHEVALLIER. *ABS: The Anti Bouncing Model for Usage Data Streams*, in "IEEE Int. Conf. on Data Mining (ICDM)", 2010, p. 1169-1174.

[11] T. M. ÖZSU, P. VALDURIEZ. *Principles of Distributed Database Systems, third edition*, Springer, 2011, http://hal.inria.fr/hal-00640392/en.

Publications of the year

Doctoral Dissertations and Habilitation Theses

[12] W. K. DEDZOE. *Traitement de Requêtes Top-k dans les Communautés Virtuelles P2P de Partage de Données*, Université de Nantes, November 2011, http://hal.inria.fr/tel-00655207/en.

[13] M. JAWAD. *Confidentialité de données dans les systèmes P2P*, Université de Nantes, June 2011, http://hal.inria.fr/tel-00638721/en.

[14] E. OGASAWARA. *Une Approche Algébrique pour les Workflows Scientifiques Orientés-Données*, Universidade Federal de Rio de Janeiro, December 2011, http://hal.inria.fr/tel-00653661/en.

[15] M. TLILI. *Infrastructure P2P pour la Réplication et la Réconciliation des Données*, Université de Nantes, June 2011, http://hal.inria.fr/tel-00643789/en.

Articles in International Peer-Reviewed Journal

[16] R. AKBARINIA, E. PACITTI, P. VALDURIEZ. *Best Position Algorithms for Efficient Top-k Query Processing*, in "Information Systems", 2011, vol. 36, n° 6, p. 973-989, http://hal.inria.fr/lirmm-00607882/en.

[17] R. AKBARINIA, M. TLILI, E. PACITTI, P. VALDURIEZ, A. A. B. LIMA. *Replication in DHTs using Dynamic Groups*, in "Journal of LNCS Transactions on Large-Scale Data and Knowledge-Centered Systems", 2011, vol. 3, p. 1-19, http://hal.inria.fr/lirmm-00607915/en.

[18] R. AKBARINIA, P. VALDURIEZ, G. VERGER. *Efficient Evaluation of SUM Queries Over Probabilistic Data*, in "IEEE Transactions on Knowledge and Data Engineering", 2012, To appear, http://hal.inria.fr/lirmm-00652293/en.

[19] A. BENOIT, H. L. BOUZIANE, Y. ROBERT. *Optimizing the reliability of streaming applications under throughput constraints*, in "International Journal of Parallel Programming", 2011, vol. 39, n° 5, p. 584-614 [*DOI :* 10.1007/s10766-011-0165-6], http://hal.inria.fr/inria-00574555/en.

[20] F. DRAIDI, E. PACITTI, B. KEMME. *A P2P Recommendation System for Large-scale Data Sharing*, in "Transactions on Large-Scale Data- and Knowledge Centered Systems", June 2011, vol. 6790, n° 3, p. 87-116, http://hal.inria.fr/lirmm-00604370/en.

[21] M. EL DICK, E. PACITTI, R. AKBARINIA, B. KEMME. *Building a Peer-to-Peer Content Distribution Network with High Performance, Scalability and Robustness*, in "Information Systems", 2011, vol. 36, n° 2, p. 222-247, http://hal.inria.fr/lirmm-00607898/en.

[22] A. MARASCU, F. MASSEGLIA. *Atypicity Detection in Data Streams: a Self-adjusting Approach*, in "Intelligent Data Analysis Journal", 2011, http://hal.inria.fr/inria-00461263/en.

[23] E. OGASAWARA, D. DE OLIVEIRA, P. VALDURIEZ, D. DIAS, F. PORTO, M. MATTOSO. *An Algebraic Approach for Data-Centric Scientific Workflows*, in "Proceedings of VLDB", 2011, vol. 4, n° 11, p. 1328-1339, http://hal.inria.fr/hal-00640431/en.

[24] F. PETITJEAN, F. MASSEGLIA, P. GANCARSKI, G. FORESTIER. *Discovering Significant Evolution Patterns from Satelllite Image Time Series*, in "International Journal of Neural Systems", 2011, vol. 21, n° 6, 15, http://hal.inria.fr/lirmm-00639480/en.

[25] J.-A. QUIANÉ-RUIZ, P. LAMARRE, P. VALDURIEZ. *Satisfaction-Based Query Replication an Automatic and Self-Adaptable Approach for Replicating Queries in the Presence of Autonomous Participants*, in "Distributed and Parallel Databases", 2011, vol. 9, n° 99, p. 1-25, To appear, http://hal.inria.fr/lirmm-00653262/en.

[26] B. SALEH, F. MASSEGLIA. *Discovering Frequent Behaviors: Time is an Essential Element of the Context*, in "Knowledge and Information Systems", 2011, vol. 28, n° 2, p. 311-331, http://hal.inria.fr/hal-00640213/en.

Invited Conferences

[27] P. VALDURIEZ. *Principles of Distributed Data Management in 2020?*, in "DEXA'11: International Conference on Databases and Expert Systems Applications", Toulouse, France, Lecture Notes in Computer Science, Springer, 2011, vol. 6860, p. 1-11 [*DOI :* 10.1007/978-3-642-23088-2], http://hal.inria.fr/hal-00640139/en.

International Conferences with Proceedings

[28] L. BARGUÑÓ, V. MUNTES-MULERO, D. DOMINGUEZ-SAL, P. VALDURIEZ. *ParallelGDB: A Parallel Graph Database Based on Cache Specialization*, in "IDEAS'11: Proceedings of the 15th Symposium on International Database Engineering & Applications", Lisbon, Portugal, B. C. DESAI, I. CRUZ, J. BERNARDINO (editors), ACM, 2011, p. 162-169 [*DOI :* 10.1145/2076623.2076643], http://hal.inria.fr/lirmm-00650603/en.

[29] W. K. DEDZOE, P. LAMARRE, R. AKBARINIA, P. VALDURIEZ. *Efficient Early Top-k Query Processing in Overloaded P2P Systems*, in "DEXA'11: International Conference on Database and Expert Systems Applications", Toulouse, France, 2011, p. 140-155, http://hal.inria.fr/lirmm-00607920/en.

[30] F. DRAIDI, E. PACITTI, M. CART, H. L. BOUZIANE. *Leveraging Social and Content-based Recommendation in P2P Systems*, in "AP2PS11: 3rd International Conference on Advances in P2P Systems", Lisbon, Portugal, November 2011, p. 13-18, http://hal.inria.fr/lirmm-00640735/en.

[31] F. DRAIDI, E. PACITTI, D. PARIGOT, G. VERGER. *P2Prec: A Social-Based P2P Recommendation System*, in "CIKM'11: Proceedings of the 20th ACM Conference on Information and Knowledge Management", United Kingdom, October 2011, p. 2593-2596, http://hal.inria.fr/lirmm-00640886/en.

[32] F. DUCHATEAU, Z. BELLAHSENE, R. COLETTA. *Matching and Alignment: What Is the Cost of User Post-Match Effort?*, in "Coopis'11 : International Conference on Cooperative Information Systems", France, 2011, p. 421-428, http://hal.inria.fr/lirmm-00643424/en.

[33] H. GOËAU, A. JOLY, S. SELMI, P. BONNET, E. MOUYSSET, L. JOYEUX. *Visual-based plant species identification from crowdsourced data*, in "MM'11 - ACM Multimedia 2011", Scottsdale, United States, ACM, November 2011 [*DOI :* 10.1145/2072298.2072472], http://hal.inria.fr/hal-00642236/en.

[34] A. JOLY, O. BUISSON. *Random Maximum Margin Hashing*, in "CVPR'11 - IEEE Computer Vision and Pattern Recognition", Colorado springs, United States, IEEE, June 2011 [*DOI :* 10.1109/CVPR.2011.5995709], http://hal.inria.fr/hal-00642178/en.

[35] P. LETESSIER, O. BUISSON, A. JOLY. *Consistent Visual Words Mining with Adaptive Sampling*, in "ICMR'11 - International Conference on Multimedia Retrieval", Trento, Italy, ACM, April 2011 [*DOI :* 10.1145/1991996.1992045], http://hal.inria.fr/hal-00642202/en.

[36] I. MAMI, R. COLETTA, Z. BELLAHSENE. *Modeling View Selection as a Constraint Satisfaction Problem*, in "DEXA'11: International Conference on Database and Expert Systems Applications", France, 2011, 15, http://hal.inria.fr/lirmm-00617288/en.

[37] D. H. NGO, Z. BELLAHSENE, R. COLETTA. *A Flexible System for Ontology Matching*, in "CAISE'11: International Conference on Advanced Information Systems Engineering", London, United Kingdom, S. NURCAN (editor), CEUR Workshop Proceedings, June 2011, p. 73-80, http://hal.inria.fr/lirmm-00639705/en.

[38] D. H. NGO, Z. BELLAHSENE, R. COLETTA. *A Generic Approach for Combining Linguistic and Context Profile Metrics in Ontology Matching*, in "ODBASE'11: International Conference on Ontologies, DataBases, and Applications of Semantics", Crete, Greece, R. MEERSMAN, T. S. DILLON, P. HERRERO, A. KUMAR, M. REICHERT, L. QING, B. C. OOI, E. DAMIANI, D. C. SCHMIDT, J. WHITE, M. HAUSWIRTH, P. HITZLER,

M. K. MOHANIA (editors), Lecture Notes in Computer Science, 2011, vol. 7045, p. 800-807, http://hal.inria.fr/lirmm-00639714/en.

[39] D. H. NGO, Z. BELLAHSENE, R. COLETTA. *YAM++ - Results for OAEI 2011*, in "OM'11: Ontology Matching Workshop at International Conference on Semantic Web", Germany, October 2011, 10, http://hal.inria.fr/lirmm-00645673/en.

[40] D. H. NGO, Z. BELLAHSENE, R. COLETTA. *YAM++ – Results for OAEI 2011*, in "ISWC'11: The 6th International Workshop on Ontology Matching", Bonn, Germany, P. SHVAIKO, J. EUZENAT, T. HEATH, C. QUIX, M. MAO, I. CRUZ (editors), CEUR Workshop Proceedings, October 2011, vol. 814, p. 228-235, http://hal.inria.fr/lirmm-00649320/en.

[41] D. PARIGOT, A. AIT LAHCEN, S. MOULINE. *Toward Data-Centric View on Service-Based Component Systems: Formalism, Analysis and Execution*, in "PDP'11: Euromicro International Conference on Parallel, Distributed and Network-Based Computing", Garching, Germany, February 2012, http://hal.inria.fr/lirmm-00648265/en.

[42] J.-A. QUIANÉ-RUIZ, P. LAMARRE, S. CAZALENS, P. VALDURIEZ. *Scaling Up Query Allocation in the Presence of Autonomous Participants*, in "DASFAA'11: International Conference on Database Systems for Advanced Applications", Hong Kong, China, Lecture notes in computer science, Springer, 2011, vol. 6588, p. 210-224 [*DOI :* 10.1007/978-3-642-20152-3_16], http://hal.inria.fr/hal-00640387/en.

[43] A. REBAI, A. JOLY, N. BOUJEMAA. *Interpretable Visual Models for Human Perception-Based Object Retrieval*, in "ICMR'11 - First ACM International Conference on Multimedia Retrieval", Trento, Italy, ACM, April 2011 [*DOI :* 10.1145/1991996.1992017], http://hal.inria.fr/hal-00642232/en.

[44] R. TRAD, A. JOLY, N. BOUJEMAA. *Large scale visual-based event matching*, in "ICMR'11 - International Conference on Multimedia Retrieval", Trento, Italy, ACM, April 2011 [*DOI :* 10.1145/1991996.1992049], http://hal.inria.fr/hal-00642210/en.

National Conferences with Proceeding

[45] F. PETITJEAN, F. MASSEGLIA, P. GANCARSKI. *Découverte de motifs d'évolution significatifs dans les séries temporelles d'images satellites*, in "EGC'11 : 11ème Conférence Internationale Francophone sur l'Extraction et la Gestion des Connaissances", Brest, France, 2011, http://hal.inria.fr/hal-00640214/en.

Conferences without Proceedings

[46] R. AKBARINIA, P. VALDURIEZ, G. VERGER. *Efficient SUM Query Processing over Uncertain Data*, in "BDA'11: Journées Bases de Données Avancées (BDA)", Rabat, Morocco, October 2011, Selected as one of the best papers, http://hal.inria.fr/hal-00639287/en.

[47] F. DRAIDI, E. PACITTI, D. PARIGOT, G. VERGER. *Demo of P2Prec: a Social-based P2P Recommendation System*, in "BDA11: 27th Journées Bases de Donnees Avancées", 2011, http://hal.inria.fr/lirmm-00640979/en.

[48] H. GOËAU, P. BONNET, A. JOLY, N. BOUJEMAA, D. BARTHÉLÉMY, J.-F. MOLINO, P. BIRNBAUM, E. MOUYSSET, M. PICARD. *The ImageCLEF 2011 plant images classi cation task*, in "ImageCLEF 2011", Amsterdam, Netherlands, September 2011, http://hal.inria.fr/hal-00642197/en.

[49] H. GOËAU, A. JOLY, I. YAHIAOUI, P. BONNET, E. MOUYSSET. *Participation of INRIA & Pl@ntNet to ImageCLEF 2011 plant images classi cation task*, in "ImageCLEF 2011", Amsterdam, Netherlands, September 2011, http://hal.inria.fr/hal-00642239/en.

[50] J.-A. QUIANÉ-RUIZ, P. LAMARRE, P. VALDURIEZ. *To Replicate or Not To Replicate Queries in the Presence of Autonomous Participants?*, in "BDA'11: Bases de Données Avancées", Rabbat, Morocco, 2011, http://hal.inria.fr/hal-00640380/en.

[51] G. VERGER, R. AKBARINIA, P. VALDURIEZ. *ProbDB: Efficient Execution of Aggregate Queries over Probabilistic Data*, in "BDA'11: Journées Bases de Données Avancées (BDA)", Rabat, Morocco, 2011, http://hal.inria.fr/hal-00639293/en.

Scientific Books (or Scientific Book chapters)

[52] Z. BELLAHSENE, A. BONIFATI, E. RAHM. *Schema Matching and Mapping*, Springer, March 2011, http://hal.inria.fr/lirmm-00581346/en.

[53] Z. BELLAHSENE, A. BONIFATI, F. DUCHATEAU, Y. VELEGRAKIS. *On Evaluating Schema Matching and Mapping*, in "Schema Matching and Mapping", Springer, March 2011, p. 253-291, http://hal.inria.fr/lirmm-00581354/en.

[54] Z. BELLAHSENE, F. DUCHATEAU. *Tuning for Schema Matching*, in "Schema Matching and Mapping", Springer, March 2011, p. 293-316, http://hal.inria.fr/lirmm-00581395/en.

[55] T. M. ÖZSU, P. VALDURIEZ. *Principles of Distributed Database Systems, third edition*, Springer, 2011, http://hal.inria.fr/hal-00640392/en.

Research Reports

[56] R. AKBARINIA, P. VALDURIEZ, G. VERGER. *SUM Query Processing over Probabilistic Data*, INRIA, May 2011, n^{o} RR-7629, http://hal.inria.fr/inria-00596020/en.

[57] A. BONIFATI, G. SUMMA, E. PACITTI, F. DRAIDI. *Semantic Query Reformulation in Social PDMS*, arXiv.org, November 2011, http://hal.inria.fr/lirmm-00640759/en.

[58] M. LIROZ-GISTAU, H. L. BOUZIANE, E. PACITTI. *Isolation Levels for Data Sharing in Large-Scale Scientific Workflows*, INRIA, May 2011, http://hal.inria.fr/lirmm-00595693/en.