EURA NOVA

Change the game with smart innovation

Master Thesis 2013 – 2014 Faculty of Science Engineering

22/03/2013

Master Thesis proposal for the academic year 2013.

TABLE OF CONTENTS

SECTION ONE

INTRODUCTION

EURA NOVA R&D

EURA NOVA is a Belgian company founded in September 2008. Our mission is simple: « Being a technologic incubator focusing on the pragmatic use of knowledge ».

Our research activities are linked to technologic directions and concrete short-term, medium-term and long-term opportunities. EURA NOVA dissociates career management from the client relationship, more in specific by basing its vision on an entrepreneurial perception of career planning.

This document presents master theses supervised by our Research and Development department. The student will work in close cooperation with the research engineers and will be invited to share his work through the in house EURA NOVA knowledge management tool.

SECTION TWO

MASTER THESIS 2013

EVOLUTION OF STORAGE SYSTEMS AND CONVERGENCE OF RELATIONAL DATABASES AND NOSQL: STUDY, MODELING AND BENCHMARK

Context: NoSQL is a large range of technologies and techniques that has been developed by the internet community to meet their scalability requirements. They made a deliberate choice to give up the transactional aspect in return of exceptional writing performances, an efficient distribution of reads and most of all storage elasticity, and this by respecting two of the three properties of the CAP theorem also known as Brewer's theorem. But today, the situation has changed and the context of the theorem has evolved [1]. The main Web actors realize that even in the case of social applications, real time and strong consistency have become crucial elements. In addition, NoSQL has been largely adopted by the community and the need to implement consistency functionalities at application level has gained importance. This is why NoSQL databases have recently evolved towards distributed architectures allowing the execution of transactions via typical leader election mechanisms. This is described by El-Abadi as *data fusion* [2]: the evolution of NoSQL structures towards RDBMS functionalities. Examples are Google Megastore [3] or G-Store [5].

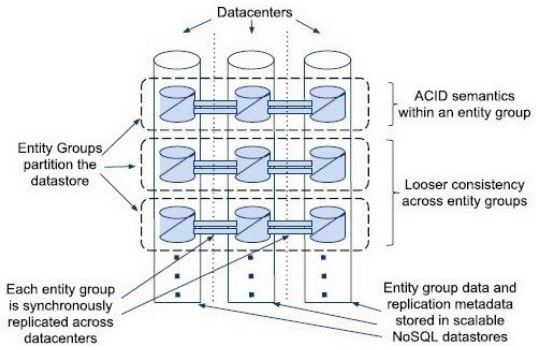

Figure 1 the Google Megastore architecture[3]

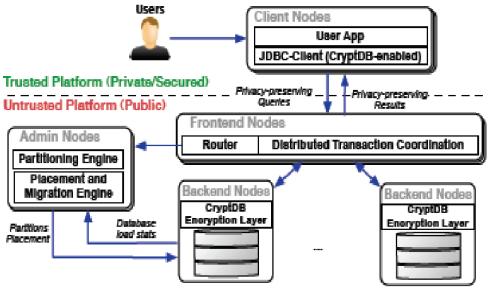

Figure 2 the MIT RelationalCloud architecture [4]

Under the pressure of NoSQL, relational databases have been gaining in performance and elasticity. The two most important evolutions are column-oriented databases and the evolution towards NoSQL functionalities. The latter has been described by El-Abadi as *data fission* [2]. Examples are Elastras, CloludSQL Server of Microsoft research or RelationalCloud of MIT.

The goal of this master thesis is to study these two major evolutions, data fission and data fusion. Besides analysing and modelling their architecture, it also aims to model their elastic behaviour, based upon research done by EURA NOVA in cooperation with the ULB and the UCL.

Contribution: the purpose of this master thesis is threefold: (1) study and model the new databases including Elastras, RelationalCloud, Spanner, Megastore and G-Store; (2) deduce their elastic nature by defining or adapting the existing models; and finally (3) validate and compare their behaviour by setting up a benchmark.

Organisation: this master thesis is organised by the ULB in collaboration with EURA NOVA R&D. The student will be supervised by the EURA NOVA R&D team.

CONVERGENCE BETWEEN HIGH PERFORMANCE GPU COMPUTING AND DISTRIBUTED DATA PROCESSING: DESIGN OF A NEW GENERATION OF COMPUTING PLATFORMS

Context: The *High Performance Computing* (HPC) community is an ecosystem that mainly consists of researchers dealing with issues related to the processing and the analysis of large data volumes. In fact, scientific computation such as used in physics, astronomy or biology, are the first real forms of Big Data. The first technologies developed within this context revolved around Grid Computing. Today the HPC community sees the emergence of computations using the graphic card processors GPU.

Figure 3 GPU have become an important source of data processing.

However, it must be pointed out that the IT industry has taken completely different directions when it comes to processing large data volumes. First of all, the internet community has invested a lot into distributed processing frameworks (Hadoop MR) and data flow Graph such as AROM [6] developed by EURA NOVA. Next, IT companies, mainly focusing on Data Warehouse, have evolved towards *in-memory* solutions by privileging *appliances* with very performing material configurations, using INFINIBAND or usNIC connections allowing RDMA reads in only microseconds of time. In fact, most of the data mining algorithms were directly *in-DB* implemented and used the maximum HW potential.

The goal of this master thesis is to validate the relevance of a high performance hybrid data processing architecture allowing the maximum use of both GPU and state-of-the-art hardware configurations, while benefiting from an architecture and a distributed processing framework to analyse data algorithms. The student will take a data mining

algorithm for classification or clusterisation as example to validate his approach. Finally, he will submit a hybrid architecture solution.

Contribution: the purpose of the master thesis is threefold: (1) study the program models and the available literature addressing the HPC GPU programming; (2) study the distributed data processing architecture; (3) submit and design a solution allowing the scheduling of a data processing on a GPU cluster deployed on several engines.

Organisation: this master thesis is organised by the ULB in collaboration with EURA NOVA R&D. The student will be supervised by the EURA NOVA R&D team.

AROM. Infinispan TERRACOTTA

STUDY OF THE USE OF A IN-MEMORY DATA GRID AS DISTRIBUTED FILESYSTEM FOR A DATA PROCESSING FRAMEWORK BASED ON DFG

Context: the distributed cache is a natural extension of the data cache concept which spans on a cluster of nodes in order to increase the transactional capacity of the cache. Recently, the latest developments in the field have led to the specification (JSR 347) of a new generation of distributed caches: the *data grids*. These extend the functionalities of traditional caches by adding indices, distribution of data and SQL, OQL-like APIs.

In the context of a past master thesis led in collaboration with the Université Libre de Bruxelles (ULB), a prototype of a distributed processing framework has been developed, named *AROM*. The framework makes use of DataFlow Graphs (DFG) for the definition of the jobs to be executed [6]. This processing model is more general and flexible compared to MapReduce and allows to define iterative and pipelined jobs in a more optimal way.

The framework currently supports the Hadoop Distributed Filesystem (HDFS) from the Apache Hadoop suite as distributed filesystem. This filesystem persists and replicates the data on the hard drive and distributes the data chunks among the computing nodes in the cluster.

The aim of this master thesis is to study and implement the support of a data grid as distributed filesystem for AROM. Furthermore the caching and flushing strategies relevant to data grids will need to be integrated in the support and take into account the types of data and job scheduled on the framework.

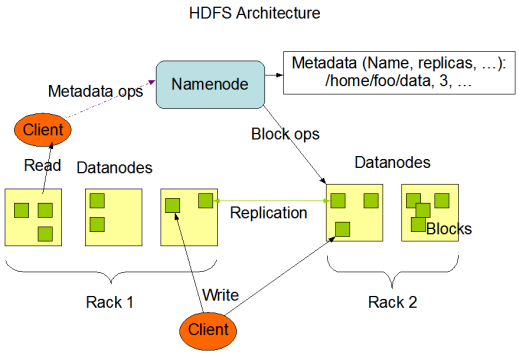

Figure 4 Architecture of HDFS

Contribution: the goal of this master thesis is (1) to study the existing data grid solutions, (2) to propose a model of the access to the data grid for the framework, (3) to develop a driver implementing the access to the data grid as filesystem for AROM according to the access model and (4) to propose data caching and flushing strategies on the data grid which are optimal with regards to the types of jobs executed on the framework.

Organisation: the master thesis is organised by the Université Catholique de Louvain (UCL) in collaboration with EURA NOVA R&D.

AROM PROCESSING FRAMEWORK – EVOLUTIONS AND NEW GENERATION OF DATA PROCESSING

Context: AROM is a distributed processing framework which aims at providing a sandbox for research and experimentation on distributed processing. The processing model is based on the general DataFlow Graphs (DFG) model which allows executing a wide range of jobs including iterative and pipelined jobs. AROM proposes a good trade-off between performance, scalability and flexibility.

The performance tests conducted in a previous study on a large scale (>40 nodes) and large amounts of data (>40 Gb) have validated the underlying concept and architecture of AROM. However these same tests have also highlighted the weaknesses of the current implementation: (1) the saturation of the asynchronous communication backbone and (2) the scheduling of the jobs on the available resources.

Concerning the first point, a first analysis has shown that the issue could be addressed by migrating the current communication stack (Akka 1.0) to a more scalable version (e.g. Akka 2.0).

On the scheduling side, the next release of Hadoop (0.23 - DotNext), the MapReduce processing framework, will aim at decoupling the management elements. To this end the scheduler (YARN) will become an independent component in the Hadoop suite and will not exclusively be tied to the MapReduce jobs. Its architecture will furthermore be modular and will provide the ability for one to define its own scheduling policies. From this point of view, YARN could represent a suitable improvement for the scheduling component of AROM.

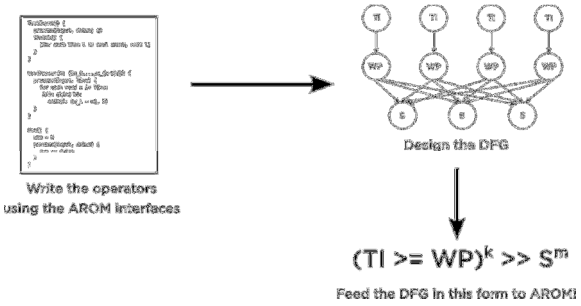

Figure 5 Definition of a AROM job

Contribution: the goal of this master thesis is (1) to propose a solution for a more scalable communication stack for AROM (e.g. Akka 2.0), (2) to implement the migration and (3) to study YARN, the modular scheduler of the Hadoop suite and propose an integration model of this scheduler to AROM

Organisation: the master thesis is organised by the Université Catholique de Louvain (UCL) in collaboration with EURA NOVA R&D.

IMPLEMENTATION OF A DIAGRAM CREATION AND MANIPULATION TOOL REPRESENTING A BUSINESS MODEL

Context: Among the rich applications, there are a lot of tools allowing to create and manage all kinds of diagrams. These diagrams are often representations of an underlying business model containing a set of data that has to be represented in one way or another (process, statistics, requests, models…).

At Eclipse level, the EMF technologies make the editing of data models quite easy. Libraries such as Graphiti provide a set of graphic representations for these models. Besides visualization, they also allow the graphical creation and manipulation of these models. This graphic editing software enables the non-technical users to manipulate these models.

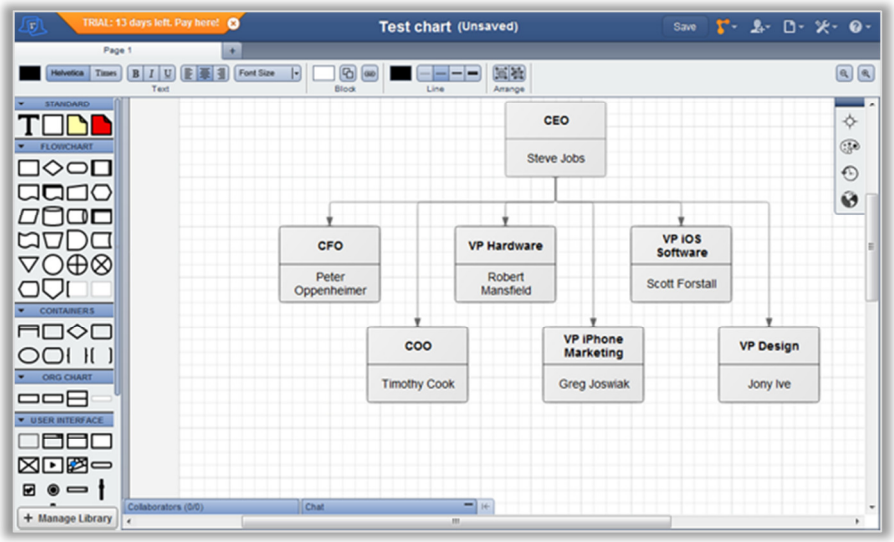

Figure 6 Most of the existing frameworks are whether owner, whether solution oriented, such as a full BPMN solution, but none of them addresses the issue of a generic modeller.

Some proprietary solutions exist but they only meet very specific needs. An equivalent tool that allows to directly edit the graphic models in a web navigator is not available. However, there are libraries enabling the visualization of data such as D3 or JIT (Javascript Infoviz Toolkit), as well as libraries able to link an underlying (potentially remotely stored) model to its visualization (e.g.: backbone.js). But you will not find a complete tool allowing the real-time editing of models.

Such a tool would enable non-technical users to elaborate models in real time without requiring any software installation on their workstation.

Contribution: This master thesis will consist of two parts:

1. The first part will study the different data representation technologies, as well as the web development environment solutions (Orion, Cloud9, …)

2. The second part will address the implementation of a tool allowing the visualisation of a remotely stored model as well as its real-time manipulation within a navigator, by using the most relevant technologies.

Organisation: this master thesis is organized by the ULB in collaboration with EURA NOVA R&D. The student will be supervised by the EURA NOVA R&D team.

PROPOSITION OF A NEW APPROACH FOR THE INDEXATION OF SOCIAL GRAPHS

Context: When using new generations of social data sources such as Facebook and Twitter, the first challenge consists in finding a structured way to represent these data, which are generally heterogeneous (digital, text, image, interaction), dynamic (evolutive) and interconnected. In addition, such a representation structure must provide easy methods to exploit and discover new knowledge. The use of graphs as a representation structure meets these requirements. Compared to usual representations such as vectors, tuples, and sets..., Graphs offer a structural representation that is characterized by a large representative capacity. Indeed, by their substructures and interactions, graphics provide a rich expressive representation through the explicit description of the data structure.

Moreover, the use of graphs as structuring models of social networks requires algorithms that allow efficient storage management and data access. So the development of index algorithms becomes a real need. Such algorithms must be designed while taking into account the three following aspect:

The structural aspect: unlike the classic index engines which only focus on the data content, the new index and search algorithms must take into account the relations and links that connect the data in the graphic. It has to be pointed out that the use of classic index graphic techniques leads to the loss of the structural data information.

The dynamic aspect: this aspect is one of the main features of internet data. The index algorithms must take into account this evolutive aspect by minimizing the laps of time needed by the search engine to collect the data and display them on the screen.

The 'scaling up' aspect: within the context of social networks and user generated content, the data generally constitute large volumes, which have become a new challenge when designing algorithms.

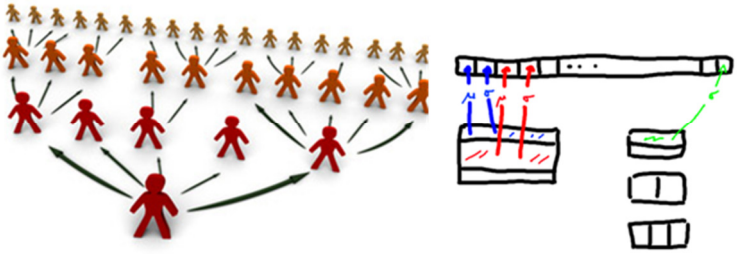

Figure 7 For the moment being, there is an important need to find new approaches to access the so called social data

Contribution: The purpose of this master thesis is threefold: (1) study the different graph index solution in general while focusing on the reachability index; (2) tailor an optimal solution for social graphs and (3) implement the solution(s) into the graph database developed by EURA NOVA.

Organisation: this master thesis is organized by the ULB in collaboration with EURA NOVA R&D. The student will be supervised by the EURA NOVA R&D team.

DESIGN OF A BIOCHEMIC *IN-MEMORY* DATA BASE FOR THE ASSISTED CONCEPTION OF MEDICINE

Context: Biochemic networks consist of a series of reactions, regulations and interactions between biochemical elements. The most important are the genetic networks that define the protein expression, the regulation of this expression through signal transduction, and the metabolic networks that are regulated by the expression of these same genes. All these interactions together constitute a system of interactions able to model the metabolism. Once these interactions are known and usable, it becomes possible to simulate the behaviour of the metabolism. This is very useful when conceiving new medicines, more specifically when simulating *in-sillico* the introduction of a molecule into the metabolism; it allows to define the activated networks and to study the potential negative effects. It also enables researchers to target the elements generating diseases or dysfunctions more efficiently by examining the global effects on the networks and help the research on certain diseases.

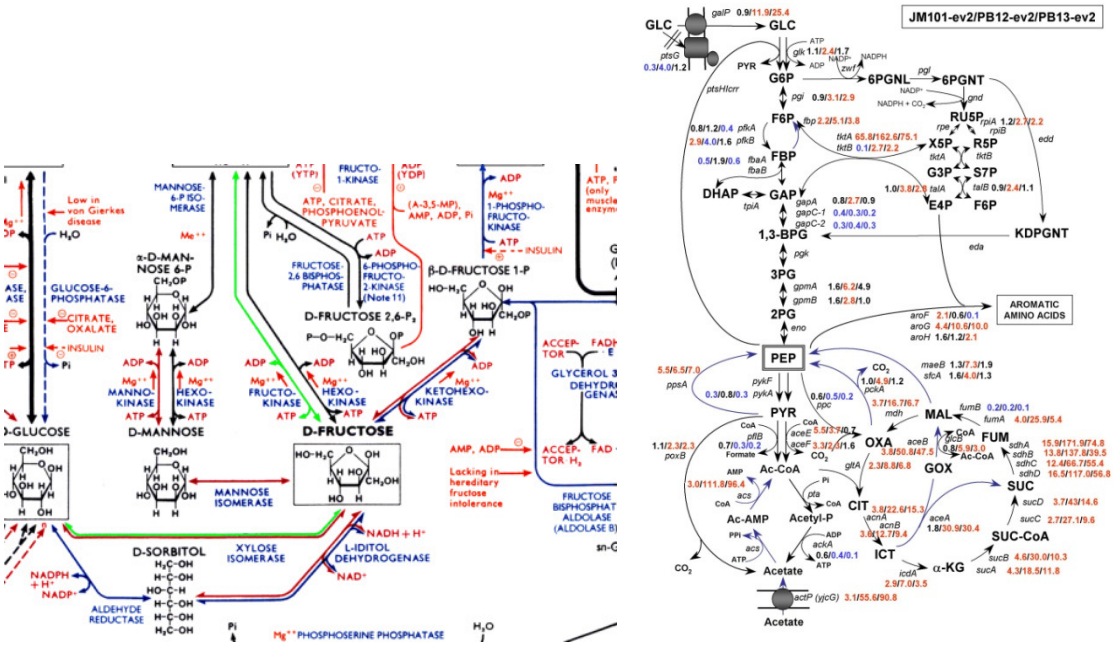

Figure 8 The biochemical networks define the interactions between the metabolic, genetic and protein networks.

Up to now this type of simulations could only represent these large graphs in two ways: in the form of files (worst option), or in relational databases which are not adapted to graph structures. Today, we can benefit from two emerging technologies: (1) the *in-memory* storage solutions such as data grids and (2) graph-oriented storage structures. In 2012-2013, EURA NOVA and the UCL have developed a first prototype of an in-memory graph database for applications allowing the analysis of social networks.

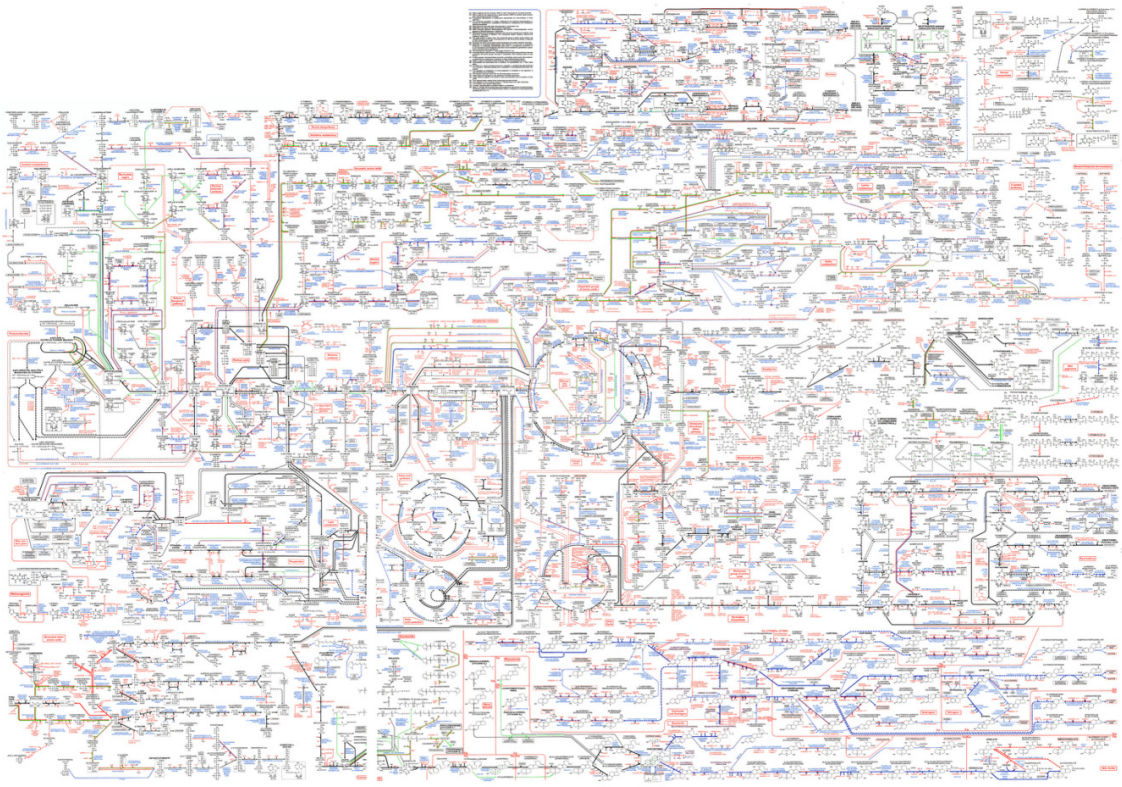

Figure 9 All the information that must be stored and analysed can be represented as a large interaction graph. The main issue to exploit this information is the graph's size and the operations to be applied [http://biochemicalpathwayswallchart.blogspot.co.uk].

The purpose of this master thesis is to adapt the technical features of this system to biochemical networks and their specific processing.

Contribution: The purpose of the master thesis is to:

1. Set up a data set op biochemical graphs based on genetic expression networks, their regulation (signal transduction) and metabolic networks (regulated by the expression of these genes). This implies the integration of graphs mainly developed by Biocyc and KEGG.

2. Define a set of extraction requests of sub-graphs which are useful for biologists and could help them for example to define the biochemical networks or the regulations activated by a molecule.

3. Deduce the type of request of the graph's *traversal*.

4. Optimize the index and the *traversal* for these requests.

Organisation: this master thesis is organised by the ULB in collaboration with EURA NOVA R&D. The student will be supervised by the EURA NOVA R&D team.

DESIGN OF A LARGE SCALE IMAGE SEARCH ENGINE

Context: Image search by content provides the ability to search images through visual queries (image query), for example when you have a picture for which you would like to obtain visually similar images. A system of image search typically consists in two parts. The first part is to describe the image. This description can be:

- Textual: in this case the image is associated with a set of words (annotations) that describe it.

- Visual: it is using recognition techniques of patterns to extract important visual characteristics of images: texture, colour, shape.

The second part is to index the image descriptors (i.e. characteristic vectors) in a way that we can quickly find similar images to a given query image based on these descriptors.

In the literature, several systems of image research have been proposed [7]. However, most of these systems are not adapted to the Big Data context. Indeed, the social network users currently publish huge amounts of images, up to terabytes per day, that require new techniques of images search by content.

Contribution: the aim of this master thesis is:

1. To study state-of-the-art techniques of images indexation.
2. To propose a suitable solution for large image databases.
3. To optimize extraction techniques of images characteristics suited for devices with low computing power (i.e. smartphone).
4. To design a distributed and elastic architecture of the system.

Organisation: this master thesis is organised by the ULB in collaboration with EURA NOVA R&D. The student will be supervised by the EURA NOVA R&D team.

METRIC LEARNING APPLIED TO THE AUTOMATIC LARGE SCALE CLASSIFICATION OF IMAGES

Context: Several techniques of supervised classification of images depend (1) on the representation of local features of the image and (2) on the metric used to calculate the similarity (or distance) between the images [8]. Recently many studies have shown the interest to learn a metric rather than use a simple metric given a priori (e.g. Euclidean distance) [9]. This approach is described in the literature as *metric learning.* Because these techniques are relatively new, there are few studies conducted in the context of classification of very large image databases.

Contribution: the aim of this master thesis is:

1. To study techniques for extracting features of images: local and global.
2. To study metric learning techniques.
3. To establish an evaluation of different approaches in the context of very large image databases.
4. To provide a distributed implementation for metric learning applied to the classification.

Organisation: this master thesis is organised by the ULB in collaboration with EURA NOVA R&D. The student will be supervised by the EURA NOVA R&D team.

PROPOSAL OF A MANAGEMENT PLATFORM FOR GRAPHS : APPLICATION TO THE MANAGEMENT OF CHEMICAL COMPONENTS

Context: The graphs are very flexible data structures that provide a high capacity for abstraction. However, the use of graphs in the context of data management (storage, indexation, search, sort, etc.) is a real challenge. The challenges to the use of graphs are met for the most basic operations such as distance calculations between two graphs. Several studies in the literature propose solutions to very specific contexts in order to facilitate the management of a set of graphs [10]. In the context of the chemical molecules management there is a gap between the scientific communities of bioinformatics and pattern recognition. This is reflected, for example, by the absence of the use of new graph matching techniques developed in the context of pattern recognition for the development of new measures of similarity between chemical components.

In addition to the measure of similarity between graphs, a solution of graphs management also requires adapted techniques of indexation.

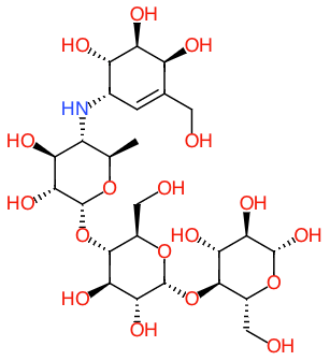

Figure 10 Chemical molecule as a graph

Contribution: the aim of this master thesis is:

1. To study matching techniques and graphs indexation.
2. To study the state-of-the-art of chemical molecules indexing.
3. To propose a platform for the management of chemical molecules represented as graphs. This platform must be designed with regards to large databases of chemical components.

Organisation: this master thesis is organised by the ULB in collaboration with EURA NOVA R&D. The student will be supervised by the EURA NOVA R&D team.

CHRONICLE MINING

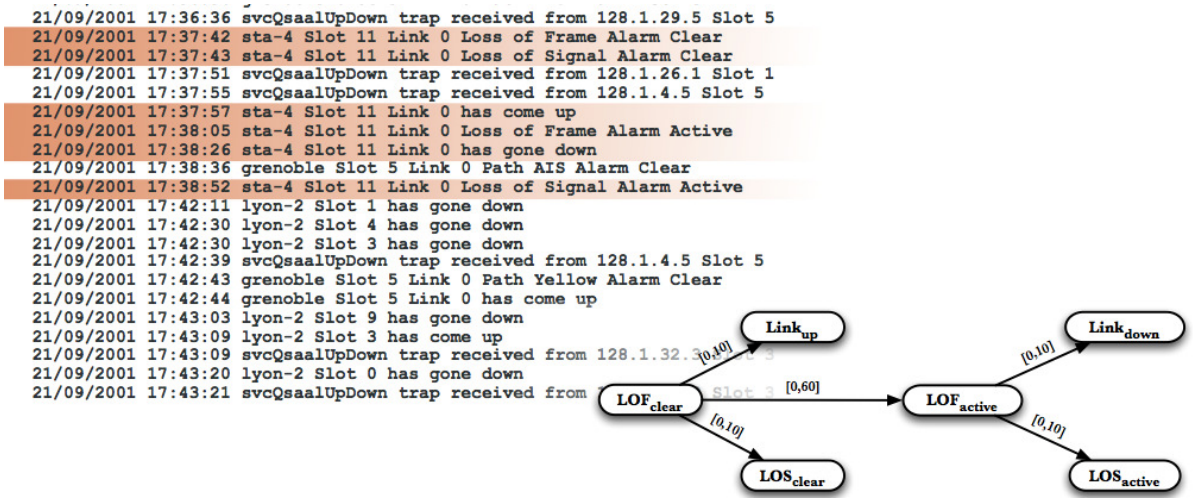

Context: A Chronicle [11, 12] is a representation of an interesting situation. In particular, it is composed of a set of timepoints separated by time intervals (timepoint 0 has to happen between 2 and 5 minutes before timepoint 1), and a set of events expected at some time points.

Chronicles are used to describe situations that a system should be monitored for; a chronicle engine takes the form of an event stream processor that takes as input raw observations on a system and emits as output the occurrences of interesting situations.

Before the system can detect a chronicle, the chronicle has to be defined. This is mostly done by experts, but should as much as possible be automated. As the computer will probably not be able to really grasp the concept of "interesting situation", we would rather aim for a collaborative tool between human expert and computer. For example, the computer would suggest a set of patterns, and the expert would then choose among them.

This opens the door to chronicle mining, which can be based on a number of existing data mining techniques, such as Petri-net based process mining [13], frequent itemset discovery [14] or trajectory mining [15].

Contribution: The goals of the thesis are :

1. Study the state of the art in Petri-nets based mining, and assess their applicability to chronicle mining.
2. Study the state of the art in frequent itemset discovery, and assess their applicability to chronicle mining.
3. Study the state of the art in trajectory mining, and assess their applicability to chronicle mining.
4. Implement a prototype of a chronicle miner.

Organisation: this master thesis is organised by the ULB in collaboration with EURA NOVA R&D. The student will be supervised by the EURA NOVA R&D team.

REFERENCES

[1] E. Brewer, *CAP twelve years later: How the "rules" have changed,* in IEEE Computer Journal vol. 45, 2012.

[2] D. Agrawal and al., *Database Scalability, Elasticity, and Autonomy in the Cloud*, in the Proceedings of the 16th international conference on Database systems for advanced applications - Volume Part I. 2011

[3] J. Baker and al., *Megastore: Providing Scalable, Highly Available Storage for Interactive Services*, in Proceedings of the Conference on Innovative Data system Research (CIDR) (2011), pp. 223-234. 2011

[4] C. Curino and al., Relational Cloud: a Database Service for the cloud, in CIDR 2011.

[5] S. Das and al., *G-Store: a scalable data store for transactional multi key access in the cloud*, in the Proceedings of the 1st ACM symposium on Cloud computing. 2010

[6] NL. Tran and al., *AROM: Processing Big Data With Data Flow Graphs and Functional Programming*, in the CRC workshop of the 4th IEEE International conference on Cloud Computing technology and Science. 2012

[7] A. W. M. Smeulders, M. Worring, S. Santini, A. Gupta, and R. Jain. Content-based image retrieval at the end of the early years. IEEE Transactions on Pattern Analysis and Machine Intelligence, 22(12):1349-1380, 2000

[8] Koen E. A. van de Sande, Theo Gevers, Cees G. M. Snoek: Evaluating Color Descriptors for Object and Scene Recognition. IEEE Transactions on Pattern Analysis and Machine Intelligence, 32(9): 1582-1596 (2010)

[9] Aurélien Bellet, Amaury Habrard, Marc Sebban: Good edit similarity learning by loss minimization. Machine Learning 89(1-2): 5-35 (2012)

[10] Salim Jouili, Salvatore Tabbone: Hypergraph-based image retrieval for graph-based representation. Pattern Recognition 45(11): 4054-4068 (2012)

[11] C. Dousson. Suivi d'évolutions et reconnaissance de chroniques. Rapport LAAS. 1994. URL: ftp://ftp.laas.fr/pub/ria/theses/christophe_dousson_16_09_94.pdf

[12] C. Dousson: Monitoring with Chronicles. (Slides) http://mklab.iti.gr/events2010/sites/default/files/2010%20-%20Events%20-%20Chronicles. pdf

[13] Toon Calders: Data mining for local patterns. EBISS 2013. URL: http://cs.ulb.ac.be/conferences/ebiss2013/files/calders_ebiss2013.pdf

[14] Wil van der Aalst: Process Mining: Making Sence of Processes Hidden in Big Event Data. EBISS 2013. URL: http://cs.ulb.ac.be/conferences/ebiss2013/files/vdaalst_ebiss2013.pdf