IN PARTNERSHIP WITH:
CNRS

Université de Bordeaux

Activity Report 2012

Project-Team CEPAGE

Algorithmics for computationally intensive applications over wide scale distributed platforms

IN COLLABORATION WITH: Laboratoire Bordelais de Recherche en Informatique (LaBRI)

RESEARCH CENTER
Bordeaux - Sud-Ouest

THEME
Distributed and High Performance Computing

Table of contents

Project-Team CEPAGE

Keywords: Distributed And Parallel Algorithms, Combinatorial Optimization, Cloud Computing, Information Indexing And Retrieval, Peer-to-peer, Scheduling

Creation of the Project-Team: January 01, 2009 .

1. Members

Research Scientists
Olivier Beaumont [Team Leader, Research Director, Inria, HdR]
Lionel Eyraud-Dubois [Research Associate (CR), Inria]
Nicolas Hanusse [Research Director (DR), CNRS, HdR]
David Ilcinkas [Research Associate (CR), CNRS]
Ralf Klasing [Research Director (DR), CNRS, HdR]
Adrian Kosowski [Research Associate (CR), Inria]

Faculty Members
Nicolas Bonichon [Assistant Professor (MdC), Université Bordeaux I, on partial Inria secondment]
Cyril Gavoille [Professor (Pr), Université Bordeaux I, IUF, HdR]
Sofian Maabout [Assistant Professor (Pr), Université Bordeaux IV, on partial Inria secondment]

External Collaborator
Philippe Duchon [Professor (Pr), Université Bordeaux I, HdR]

Engineer
Lucas Verdonk [Engineer, Euler European grant (October 2011 –)]

PhD Students
Florent Foucaud [Université Bordeaux I, MENRT grant, (since September 2009)]
Przemyslaw Uznanski [Inria / CRA (since October 2010)]
Halftermeyer Pierre [MESR grant (since October 2010)]
Quentin Godfroy [University Bordeaux I, ANR Aladdin grant (since October 2009)]
Christian Glacet [Inria / FP7 EULER Grant (since November 2010)]
Ahmed Mouhamadou Wade [Université Bordeaux I, MENRT grant (since February 2011)]
Dominik Pajak [Inria (since October 2011)]

Administrative Assistants
Marie Sanchez [Inria]
Nicolas Jahier [Inria]

2. Overall Objectives

2.1. General Objectives

CEPAGE project was launched (as a team) in 2006, when the development of interconnection networks had led to the emergence of new types of computing platforms, characterized by the heterogeneity of both their processing and communication resources, their geographical dispersion, the absence of centralized control, and their instability in terms of the number and performance of participating resources. At the same time, P2P file exchange systems and Volunteer Computing platforms were developing, reinforcing the idea that big platforms could be built by aggregating a large set of small distributed resources. In the context of intensive computations, the solutions were less sophisticated than in the context of file exchanges, and mostly relied on the client-server model (SETI@home, folding@home), with centralized knowledge and control of the platform, and no direct communications between peers. In this context, the general goal of CEPAGE was to come up with new strategies in order to extend the set of computation-intensive applications that could be run on large scale distributed platforms, by gathering researchers with expertize in scheduling of tasks and collective communications, graph theory, design of overlay networks, modeling of network topologies, small world networks, distributed algorithms, compact data structures, routing and randomized algorithms, later extended to include mobile agents and databases.

2.1.1. Objectives for the evaluation period

The above context lead us to identify the following research axes (extracted from the project proposal).

1. Models
   1. At a low level, to understand the underlying physical topology and to obtain realistic models whose parameters can be instantiated at runtime.
   2. At a higher level, to derive models of the dynamism of targeted platforms, both in terms of participating resources and resource performance.
2. Overlays and distributed algorithms:
   1. To understand how to augment the logical topology in order to achieve good properties of P2P systems.
   2. To build overlays dedicated to specific applications and services that achieve good performance.
   3. To understand how to dynamically adapt scheduling algorithms (in particular collective communication schemes) to changes in network performance and topology, using randomized algorithms.
3. Compact and distributed data structures:
   1. To understand how to dynamically adapt compact data structures to changes in network performance and topology.
   2. To design sophisticated labeling schemes in order to answer complex predicates using local labels only.

Regarding the 2007-2011 objectives of the projects (1)-(3), the vast majority of them has been fulfilled as shown in Sections 4.1, 4.2 and 4.3.

Nevertheless, during the last four years, several adjustments were made with respect to the initial objectives.

- First, we under-evaluated in our project proposal the importance of issues related to security and anonymity, by saying that we would rely on external mechanisms provided by distributed systems to ensure them. It turned out that the question was harder to solve and a necessary condition in order to deploy our algorithms at a large scale. Therefore, we spent a lot of effort to work on these issues using the behavior and capacities of mobile agents under adversarial models. The recruitments of David Ilcinkas (CR CNRS, 2007) and Adrian Kosowski (CR Inria, 2010) have strongly reinforced this axis.
- The picture changed due to the coming of Cloud Computing as a mean to "renting" external resources rather than aggregating and sharing privately-owned ones. This solution is particularly appealing in the context of computing-intensive applications. Our efforts to develop widely used distributed codes for molecular dynamics simulations and continuous integration were made unsuccessful by this new context. We are moving toward this direction, in particular by considering stability of resource allocation problems, reliability, and stochastic optimization issues.

2.2. Goal and Context

2.2.1. General Context

The first set of questions that we consider is related to distributed computations, where the Internet is the underlying network. Since the topology of the underlying network is unknown, the use of logical networks (overlays) is required. In turn, the choice of the overlay will have an impact on the complexity of the algorithms. In this context, only the performance of the whole chain is meaningful, which requires to collect raw data and then to propose network models and algorithms based on these models, such that the performance of the resulting algorithm is good on raw data. This also requires studying the influence of the topology of the overlay network (the underlying graph) on the complexity of fundamental questions, such as graph exploration or black-hole search.

The second set of questions is related to distributed data structures. In general, the question is related to the compromise between the size of the data structure to be stored on each node and the time to answer a request (estimate the bandwidth between two nodes, compute the closest common ancestor of two nodes in a tree) or to perform a task (route a message in a network based on the information stored at the router nodes).

In order to study these questions, our research plan is based on the following goals. Firstly, we aim both at building strong foundations for distributed algorithms (graph exploration, black-hole search,...) and distributed data structures (routing, efficient query, compact labeling...) to understand how to explore large scale networks in the context of failures and how to disseminate data so as to answer quickly to specific queries. Secondly, we aim at building simple (based on local estimations without centralized knowledge), realistic models to accurately represent resource performance and to build a realistic view of the topology of the network (based on network coordinates, geometric spanners, δ-hyperbolic spaces). Then, we aim at proving that these models are tractable by providing low complexity distributed and randomized approximation algorithms for a set of basic scheduling problems (independent tasks scheduling, broadcasting, data dissemination,...) and associated overlay networks. At last, our goal is to prove the validity of our approach through softwares dedicated to several applications (molecular dynamics simulations, continuous integration) as well as more general tools related to the model we propose (AlNEM for automatic topology discovery, SimGRID for simulations at large scale) and collections of datasets (Hubble for the continuous integration DAGs, Bedibe for latency and bandwidth measurements).

For the sake of the clarity of the presentation of our contributions during the last evaluation period, we decided to perform a synthesis, organizing our work into three main axes.

1. Resource Allocation and Scheduling (Section 4.1), with an emphasis on the interaction between network performance modeling and the design of efficient and guaranteed algorithms (covers axes 1(a), 2(a) and 2(b) of Section 2.1.1);
2. Compact Routing (Section 4.2), with an emphasis on the design of specific strategies for restricted graph classes and the design of data structures resilient to resource failures (covers axes 1(b), 2(b), 2(c), 3(a), 3(b) of Section 2.1.1);
3. Mobile Agents (Section 4.3), with en emphasis on the detection of dynamic faults and improvement of dissemination of information (covers axes 1(b), 2(b), 3(b) of Section 2.1.1);.

Nevertheless, strong relationships and collaborations exist between these axes:

- the design and analysis of overlay networks is shared by Resource Allocation and Scheduling and Mobile Agents axes (see for example [61], [85], [54], [101], [77]),
- the design of geometric spanners and geographic routing protocols is shared by Resource Allocation and Scheduling and Compact Routing axes (see for example [85], [69], [70], [105]),
- the use of specific graph classes (bounded tree-width, ...) is shared by Mobile Agents and Compact Routing axes (see for example [93], [100], [82], [87], [84], [46]).

2.3. Highlights of the Year

- A long paper has been presented at the STOC'12 Conference: I. Abraham, S. Chechik and C. Gavoille. *Fully Dynamic Approximate Distance Oracles for Planar Graphs via Forbidden-Set Distance Labels* in 44th Annual ACM Symposium on Theory of Computing (STOC), pp. 1199-1217, New York, May 2012.
- Ralf Klasing was the Conference Chair of the *11th International Symposium on Experimental Algorithms (SEA 2012)*, Bordeaux, France, June 7-9, 2012.
- The members of CEPAGE have been involved in the following program committees: SODA 2013, IPDPS 2013, DISC 2012, ISAAC 2012, ICDCN 2012, IPDPS 2012 FOMC 2012, ADHOC-NOW 2012, IWOCA 2012, SEA 2012, ALGOTEL 2012, FUN2012.

3. Scientific Foundations

3.1. Modeling Platform Dynamics

Modeling the platform dynamics in a satisfying manner, in order to design and analyze efficient algorithms, is a major challenge. In distributed platforms, the performance of individual nodes (be they computing or communication resources) will fluctuate; in a fully dynamic platform, the set of available nodes will also change over time, and algorithms must take these changes into account if they are to be efficient.

There are basically two ways one can model such evolution: one can use a *stochastic process*, or some kind of *adversary model*.

In a stochastic model, the platform evolution is governed by some specific probability distribution. One obvious advantage of such a model is that it can be simulated and, in many well-studied cases, analyzed in detail. The two main disadvantages are that it can be hard to determine how much of the resulting algorithm performance comes from the specifics of the evolution process, and that estimating how realistic a given model is – none of the current project participants are metrology experts.

In an adversary model, it is assumed that these unpredictable changes are under the control of an adversary whose goal is to interfere with the algorithms efficiency. Major assumptions on the system's behavior can be included in the form of restrictions on what this adversary can do (like maintaining such or such level of connectivity). Such models are typically more general than stochastic models, in that many stochastic models can be seen as a probabilistic specialization of a nondeterministic model (at least for bounded time intervals, and up to negligible probabilities of adopting "forbidden" behaviors).

Since we aim at proving guaranteed performance for our algorithms, we want to concentrate on suitably restricted adversary models. The main challenge in this direction is thus to describe sets of restricted behaviors that both capture realistic situations and make it possible to prove such guarantees.

3.2. Models for Platform Topology and Parameter Estimation

On the other hand, in order to establish complexity and approximation results, we also need to rely on a precise theoretical model of the targeted platforms.

- At a lower level, several models have been proposed to describe interference between several simultaneous communications. In the 1-port model, a node cannot simultaneously send to (and/or receive from) more than one node. Most of the "steady state" scheduling results have been obtained using this model. On the other hand, some authors propose to model incoming and outgoing communication from a node using fictitious incoming and outgoing links, whose bandwidths are fixed. The main advantage of this model, although it might be slightly less accurate, is that it does not require strong synchronization and that many scheduling problems can be expressed as multi-commodity flow problems, for which efficient decentralized algorithms are known. Another important issue is to model the bandwidth actually allocated to each communication when several communications compete for the same long-distance link.

- At a higher level, proving good approximation ratios on general graphs may be too difficult, and it has been observed that actual platforms often exhibit a simple structure. For instance, many real life networks satisfy small-world properties, and it has been proved, for instance, that greedy routing protocols on small world networks achieve good performance. It is therefore of interest to prove that logical (given by the interactions between hosts) and physical platforms (given by the network links) exhibit some structure in order to derive efficient algorithms.

3.3. General Framework for Validation

3.3.1. Low level modeling of communications

In the context of large scale dynamic platforms, it is unrealistic to determine precisely the actual topology and the contention of the underlying network at application level. Indeed, existing tools such as Alnem [103] are very much based on quasi-exhaustive determination of interferences, and it takes several days to determine the actual topology of a platform made up of a few tens of nodes. Given the dynamism of the platforms we target, we need to rely on less sophisticated models, whose parameters can be evaluated at runtime.

Therefore, we propose to model each node using a small set of parameters. This is related to the theoretical notion of distance labeling [92], and corresponds to assigning labels to the nodes, so that a cheap operation on the labels of two nodes provides an estimation of the value of a given parameter (the latency or the bandwidth between two nodes, for instance). Several solutions for performance estimation on the Internet are based on this notion, under the terminology of Network Coordinate Systems. Vivaldi [83], IDES [104] and Sequoia [106] are examples of such systems for latency estimation. In the case of bandwidth estimation, fewer solutions have been proposed. We have studied the last-mile model, in which we model each node by an incoming and an outgoing bandwidth and neglect interference that appears at the core of the network (Internet), in order to concentrate on local constraints.

3.3.2. Simulation

Once low level modeling has been obtained, it is crucial to be able to test the proposed algorithms. To do this, we will first rely on simulation rather than direct experimentation. Indeed, in order to be able to compare heuristics, it is necessary to execute those heuristics on the same platform. In particular, all changes in the topology or in the resource performance should occur at the same time during the execution of the different heuristics. In order to be able to replicate the same scenario several times, we need to rely on simulations. Moreover, a metric for providing approximation results in the case of dynamic platforms necessarily requires computing the optimal solution at each time step, which can be done off-line if all traces for the different resources are stored. Using simulation rather than experiments can be justified if the simulator itself has been proven valid. Moreover, the modeling of communications, processing and their interactions may be much more complex in the simulator than in the model used to provide a theoretical approximation ratio, such as in SimGrid. In particular, sophisticated TCP models for bandwidth sharing have been implemented in SimGRID.

During the course of the USS-SimGrid ANR Arpege project, the SimGrid simulation framework has been adapted to large scale environments. Thanks to hierarchical platform description, to simpler and more scalable network models, and to the possibility to distribute the simulation of several nodes, it is now possible to perform simulations of very large platforms (of the order of 10^5 resources). This work will be continued in the ANR SONGS project, which aims at making SimGrid usable for Next Generation Systems (P2P, Grids, Clouds, HPC). In this context, simulation of exascale systems are envisioned, and we plan to develop models for platform dynamicity to allow realistic and reproducible experimentation of our algorithms.

3.3.3. Practical validation and scaling

Finally, we propose several applications that will be described in detail in Section 5. These applications cover a large set of fields (molecular dynamics, continuous integration...). All these applications will be developed and tested with an academic or industrial partner. In all these collaborations, our goal is to prove that the services that we propose can be integrated as steering tools in already developed software. Our goal is to assert the practical interest of the services we develop and then to integrate and to distribute them as a library for large scale computing.

At a lower level, in order to validate the models we propose, i.e. make sure that the predictions given by the model are close enough to the actual values, we need realistic datasets of network performance on large scale distributed platforms. Latency measurements are easiest to perform, and several datasets are available to researchers and serve as benchmarks to the community. Bandwidth datasets are more difficult to obtain, because of the measurement cost. As part of the bedibe software (see section 5.4), we have implemented a script to perform such measurements on the Planet-Lab platform [72]. We plan to make these datasets available to the community so that they can be used as benchmarks to compare the different solutions proposed.

4. Application Domains

4.1. Resource Allocation and Scheduling

4.1.1. Project-team positioning

CEPAGE has undertaken tasks related to the *high level modeling* of heterogeneous networks, both at logical level (overlay networks design) and performance level (latency, bandwidth prediction, connectivity artifacts) in order to optimize tasks such as *resource allocation* and *scheduling* of computations and communications. Objectives include replica placement, broadcasting (streaming) of large messages, independent tasks scheduling and optimization of OLAP databases. Such problems have received at lot of attention in research centers in the USA (Armherst, Colorado, ...), in Spain (Madrid), Poland (Wroclaw), Germany (Dortmund), and others. Papers on algorithmic aspects of platform modeling, scheduling and resource allocation appear at parallel processing conferences and journals in Parallel and Distributed Computing (IPDPS, EuroPar, HIPC, SPAA, IEEE TPDS, JPDC) and members of CEPAGE are strongly involved in many of these events (IPDPS, EuroPar, TPDS) as well as helping to animate well-established specialized workshops, such as HCW and HeteroPar.

Within Inria, studies on overlay networks are performed in the ASAP and GANG projects, and studies related to scheduling and resource allocation are done within the ROMA and the MOAIS projects (and to some extent within ALGORILLE).

4.1.2. Scientific achievements

The approach followed in the CEPAGE project, and our main originality, is to consider the whole chain, from gathering actual data on the networks to platform modeling and complexity analysis. Indeed, many complexity analysis studies are performed on models whose parameters cannot actually be evaluated (this applies, for instance, to all algorithms that assume that the topology of a platform running over the Internet is known in advance) and many platform models are intractable from an algorithmic perspective (this applies, for instance, to all models that represent latencies or bandwidths between all pairs of nodes as a general matrix). Our general goal is to provide models whose parameters can be evaluated at runtime using actual direct measurements, to propose algorithms whose worst-case (or average-case) behavior can be proved for this model, and finally to evaluate the whole chain (model + algorithm + implementation).

From an applicative perspective, in the framework of the PhD Thesis of Hejer Rejeb, we have considered several storage and resource allocation problems in collaboration with Cyril Banino-Rokkones at Yahoo! Trondheim (dealing with actual datasets enabled us to improve known approximation results in this specific context). We have in particular studied the modeling of TCP mechanism for handling contentions and its influence on the performance of several scheduling algorithms and advocated the use of QoS mechanisms for prescribed bandwidth sharing (IPDPS 2010 [67], ICPADS 2008 [52], AlgoTel 2009 [64], ICPADS 2009 [63], PDP 2010 [65]). In the PhD thesis of Hubert Larchevêque, we have considered the problem of aggregating resources (or placing replicas) in a distributed network (Sirocco 2008 [54], Opodis 2008 [55], ICPP 2011 [60], AlgoTel 2011 [56]) so that each group satifies some properties (in terms of aggregated memory, CPU and maximal distance in terms of latency within a group). We proved several multi-criteria approximation results for this problem, and we compared several embedding tools (Vivaldi, Sequoia) in the context of resource aggregation. For these applications, we have also provided when possible distributed algorithms based on sophisticated overlay networks, in particular in order to deal with heterogeneity (IPDPS 2008 [61]). In the PhD Thesis of Przemyslaw Uznanski, we focus on the design of efficient streaming and broadcasting strategies, in particular in presence of connectivity artifacts like firewalls (IPDPS 2010 [62], ICPADS 2011 [59]). We have also worked on establishing under the bounded multiport model several new complexity results for classical distributed computing models such as divisible load theory (HCW 2008 [57], IPDPS 2008 [107], IPDPS 2012 [58]) that have been later extended to Continuous Integration (HCW 2012 [53]).

In the context of database query optimization, materializing some queries results for optimization is a standard solution when execution time performance is crucial. In the datacube context, the problem has been studied for a long time under the storage space limit constraint. Here also, we were able to reformulate this problem by considering instead the execution time as the hard constraint while the objective is to reduce the storage space. Even if the problem turns to be NP-hard, this reformulation allowed us to provide effective approximate solutions with both space and performance bounded guarantees (EDBT 2009 [96]). Moreover, reducing the storage space tends to reduce the maintenance time since the latter is linearly proportional to the former. Finally, we characterized the minimal number of updates to be performed before performance becomes no more guaranteed and a new solution must be recomputed (ADBIS 2008 [97]). One of the key concepts we used for solving this problem was that of a *border*. It turns out that this notion is equivalent to e.g., maximal frequent itemsets or minimal functional dependencies extensively studied by data mining community. In contrast to all previous proposals, we proposed the only parallel algorithm computing these borders with a speed-up guarantee regarding the number of processing units (CIKM 2011 [95]). Besides the analytical study, its implementation in maximal frequent itemset mining outperforms state of the art implementations (see Section 5.1).

To achieve these results, our efforts have also focused on analyzing and building realistic datasets (AlgoTel 2012 [86]) and proposing data analysis results for specific distributions (ISAAC 2011 [48]). On the modeling side, in general, for bandwidth and contention modeling, we have proved that the bounded multi-port model (where each node is associated to an incoming bandwidth, an outgoing bandwidth and a maximal number of simultaneous TCP connexions) is both implementable, realistic and tractable (EuroPar 2011 [66]). In particular, we have proved in strongly different contexts (allocation of virtual machines to physical machines, overlay design for broadcasting, server allocation for volunteer computing) that the use of resource augmentation enables to obtain quasi-optimal results. All our modeling efforts and algorithms have been included into the SimGRID Software (http://simgrid.gforge.inria.fr), which enables us both to compare several algorithms under the same exact conditions and to compare the results obtained with several communication models (see Section 5.1)..

Perspectives: We believe that our approach based on sound models, approximation algorithms for these models, followed by experimental validation is a strong one and we intend to continue in this direction in the following years. Our goal of designing realistic solutions pushes towards considering average case analysis of our algorithms, as well as robust optimization techniques. Furthermore, the recent strong interest in Cloud systems from the community entices us to use our expertise in resource allocation for the optimization of Cloud systems, both from the provider and from the user points of view. We already have some interesting contacts with local companies to share start collaborating on these topics. In this context, reliability issues are very important, and we believe that robust optimization is a very relevant approach for these problems.

4.2. Compact Routing

4.2.1. Project-team positioning

In this axis, CEPAGE mainly works on the design on distributed and light data structures. One of the techniques consists in summarizing the topology and metric of the networks allowing to route or to approximate the original distances within the network. Such structures, often called *spanners*, does not require the storage of all the original network links. Then we get economic distributed data structures that can be updated without a high communication cost. Our main collaborations are done with the best specialists world-wide, in particular: Israel (Weizmann), USA (MIT, Microsoft, Chicago), Belgium (Alcatel Lucent-Bell), France (Paris, Nice).

Algorithms and Routing are also intensively studied in research labs in the USA (CAIDA). Our contributions appear regularly at all of the major conferences in Distributed Computing (PODC, DISC, SPAA), as well as at top venues with a more general algorithmic audience (STOC, SODA, ICALP, ESA). Members of CEPAGE actively participate in these events (ICALP 2010 and DISC 2009 were organized by members of CEPAGE).

Within Inria, studies of mobile agents are also performed in the GANG project and to some extent also within MASCOTTE within the european project EULER.

4.2.2. Scientific achievements

There are several techniques to manage sub-linear size routing tables (in the number of nodes of the platform) while guaranteeing almost shortest paths. Some techniques provide routes of length at most $1+\epsilon$ times the length of the shortest one while maintaining a poly-logarithmic number of entries per routing table. However, these techniques are not universal in the sense that they apply only on some class of underlying topologies. Universal schemes exist. Typically they achieve $O(\sqrt{n})$-entry local routing tables for a stretch factor of 3 in the worst case. Some experiments have shown that such methods, although universal, work very well in practice, in average, on realistic scale-free or existing topologies.

The space lower bound of $O(\sqrt{n})$-entry for routing with *multiplicative* stretch 3 is due to the existence of dense graphs with large girth. Dense graphs can be sparsified to subgraphs (spanners), with various stretch guarantees. There are spanners with *additive* stretch guarantees (some even have constant additive stretch) but only very few additive routing schemes are known.

In (SPAA 2012 [90]), we give reasons why routing in unweighted graphs with *additive* stretch is difficult in the form of space lower bounds for general graphs and for planar graphs. On the positive side, we give an almost tight upper bound: we present the first non-trivial compact routing scheme with $o(\lg^2 n)$-bit addresses, *additive* stretch $O(n^{1/3})$, and table size $O(n^{1/3})$ bits for planar graphs.

We have recently considered the *forbidden-set* extension of distance oracles and routing schemes. Given an arbitrary set of edge/node failure F, a source s and a target t such that $s,t \notin F$, the goal is to route (or evaluate the distance) between s and t in the graph $G \setminus F$, so avoiding F. The classical problem is for $F = \varnothing$. This extension is considered as a first step toward fully dynamic data-structures, a challenging goal. For graphs of low doubling dimension we have shown in (PODC 2012 [47]) that it is possible to route from s to t in $G \setminus F$ with stretch $1+\epsilon$, for all s,t,F, given poly-logarithmic size labels of all the nodes invoked in the query (s,t,F). This has been generalized to all planar graphs achieving similar stretch and label size performences. As a byproduct we have designed a fully dynamic algorithm for maintaining $1+\epsilon$ approximate distances in planar graphs supporting edge/node addition/removal within update and query time $\sqrt{n}$ in the worst-case (STOC 2012 [46]).

Θ_k-graphs are geometric graphs that appear in the context of graph navigation. The shortest-path metric of these graphs is known to approximate the Euclidean complete graph up to a factor depending on the cone number k and the dimension of the space. We have introduced in (WG 2010 [68]) a specific subgraph of the Θ_6-graph defined in the 2D Euclidean space, namely the half-Θ_6-graph, composed of the even-cone edges of the Θ_6-graph. Our main contribution is to show that these graphs are exactly the TD-Delaunay graphs, and are strongly connected to the geodesic embeddings of orthogonal surfaces of coplanar points in the 3D Euclidean space. We also studied the asymptotic behavior of these spanners (*Adv. in Appl. Proba.* [105]) and in collaboration with Ljubomir Perković, we worked on the question of bounded degree planar spanner. We proposed an algorithm that computes a plane 6-spanner of degree at most 6 in (ICALP 2010 [69]). The previous best bound on the maximum degree for constant stretch plane spanners was only 14.

In order to cope with network dynamism and failures, and motivated by multipath routing, we introduce a multi-connected variant of spanners. For that purpose we introduce in (OPODIS 2011 [91]) the p-multipath cost between two nodes u and v as the minimum weight of a collection of p internally vertex-disjoint paths between u and v. Given a weighted graph G, a subgraph H is a p-multipath s-spanner if for all u,v, the p-multipath cost between u and v in H is at most s times the p-multipath cost in G. The s factor is called the stretch. Building upon recent results on fault-tolerant spanners, we show how to build p-multipath spanners of constant stretch and of $O(n^{1+1/k})$ edges, for fixed parameters p and k, n being the number of nodes of the graph. Such spanners can be constructed by a distributed algorithm running in $O(k)$ rounds. Additionally, we give an improved construction for the case $p = k = 2$. Our spanner H has $O(n^{3/2})$ edges and the p-multipath cost in H between any two node is at most twice the corresponding one in G plus $O(W)$, W being the maximum edge weight.

We also worked on compact coding in data warehouses: in order to get quick answer in large data, we have to estimate, select and materialize (store) partial data structures. We got several solutions with a prescribed guarantee in different models for the following problems: view size estimation with small samples, view selection, parallel computation of frequent itemsets. In (*Theor. Comp. Sci.* [94]) a new algorithm that allow the administrator or user of a DBMS to choose which part of the data cube to optimize (known as the *the views selection problem*), that takes as input a fact table and computes a set of views to store in order to speed up queries.

Perspectives: The compact coding activity in data-warehouse is promising since the amount of data collected keeps on increasing and being able to answer in real-time complex requests (data mining) is still challenging.

Some robust data structures already exist which, given a small number of k changes of topology or k faults, tolerate these faults, i.e., alternative routes with bounded stretch can be provided without any updates. This is a first step toward dynamic networks but the updates of these data structures are currently still quite complicated with a high communication cost.

4.3. Mobile Agents

4.3.1. Project-team positioning

CEPAGE has undertaken tasks related to the design of algorithms which control the behavior of so called *mobile agents*, moving around a network or a geometric environment, with the goal of achieving a specified objective. Objectives of central importance to the study include: exploration of unknown environments, terrain patrolling, network maintenance, and coordination of activities with other agents. Such problems have in recent years been the object of interest of numerous research teams working on Distributed Computing worldwide, in particular, at research centers in Canada (Quebec), Israel (Tel Aviv, Haifa), France (Paris, Marseille), the UK (London, Liverpool), and Switzerland (Zurich). Algorithms for mobile agents in social networking applications are also intensively studied in research labs in the USA (Stanford, Facebook). Papers on mobile agents appear regularly at all of the major conferences in Distributed Computing (PODC, DISC, SPAA), as well as at top venues with a more general algorithmic audience (SODA, ICALP, ESA). Members of CEPAGE actively participate in these events, and are also a recognizable part of the European community focused around mobile agents, helping to animate well-established specialized conferences, such as SIROCCO and OPODIS.

Within Inria, studies of mobile agents are also performed in the GANG project, and to some extent also within MASCOTTE. CEPAGE has active research links with both of these teams.

4.3.2. Scientific achievements

The work of CEPAGE has focused on contributing new decentralized algorithms for controlling mobile entities known as *agents*, deployed in unknown environments. We mainly considered the network setting, in which agents moving around the nodes of the network graph may be used to analyze the structure of the network and to perform maintenance tasks, such as detecting dynamic faults, improving/monitoring dissemination of information, etc. Our theoretical studies focused on designing new strategies for controlling the behavior of agents and answering crucial questions concerning the feasibility of solving fundamental problems, subject to different model assumptions and constraints on the knowledge and computational power of agents.

One major line of our research focused on the so called *anonymous graph model* in which an agent is unable to determine the identifier of the node of its current location, but can only see a local ordering of the links around it. Such a study is motivated e.g. by scenarios in which the identifiers of nodes may be too large for the agent to process using its bounded resources, or may change in time. In this model, we studied two of the most fundamental problems: that of traversing all of the nodes of the network (exploration) and of meeting another agent in the network (rendezvous), so as to coordinate with it. Our contributions include a precise characterization of the space requirements for agents solving both of these problems deterministically: exploration in (*Trans. Alg.* 2008 [73]) and rendezvous in (*Dist. Comp.* 2012 [81]), in a paper presented at the Best Paper Session of PODC 2010. We have also studied fast solutions for specific scenarios of the rendezvous problem (DISC 2010 [49], DISC 2011 [74], SPAA 2012 [82]) and the problem of approximate map construction within an anonymous graph (OPODIS 2010 [71]). A separate problem, intensively studied in recent years by several research teams, concerns the exploration of a network with pre-configured ports so as to assist the agent. In our work on the topic, our team has proposed several new techniques for graph decomposition, leading in particular to the shortest currently known strategies of periodic exploration for both the case of memoryless (*Algorithmica* 2012 [101]) and small-memory agents (SIROCCO 2009 [77]).

A closely related line of research was devoted to the design of network exploration strategies which guarantee a fast and fair traversal of all the nodes, making use of agents with extremely restricted capabilities. Such strategies were inspired by the random walk, but had the additional advantage of deterministic and desirable behavior in worst-case scenarios. We presented a series of results in the area at notable conferences, involving both the design of new exploration strategies (ICALP 2009 [75]) and completely new insights into previously known approaches such as the so called "rotor-router model" (DISC 2009 [50], OPODIS 2009 [51]). All of the proposed algorithms were shown to be viable alternatives to the random walk, competing in terms of such parameters as cover time, steady-state exploration frequency, and stabilization in the event of faults.

Our efforts have also focused on the theory of coordinating activities of large groups of agents. We have conducted pioneering work in the so called look-compute-move model in networks, in which extremely restricted (asynchronous and oblivious) agents, relying on snapshot views of the system, are nevertheless able to perform useful computational tasks. Our solutions to the problems of collective exploration in trees (*Theor. Comp. Sci.* 2010 [88]) and gathering agents on a ring (*Theor. Comp. Sci.* 2008 [99] and 2010 [98]) have sparked a long line of follow-up research, accumulating more than 120 citations in total (according to Google Scholar). In a slightly different scenario, we have considered computations with teams of agents whose task is to collaboratively detect and mark potentially dangerous (faulty) links of the network, called "black holes", which are capable of destroying agents which enter them. We have provided important contributions to the theory of black hole search in both undirected (SIROCCO 2008 [76], DISC 2008 [89]) and directed (*Theor. Comp. Sci.* [102]) graphs.

It is expected that the mobile agent theme of CEPAGE will give rise to 2 PhD theses. In 2013, Ahmed Wade will defend his thesis on mobile agent protocols for dynamic networks, whereas in 2014 Dominik Pajak will defend his thesis on multi-agent protocols for efficient graph exploration. Our scientific interests also include mobile agent protocols for geometric applications, more remote from the central themes of CEPAGE, but having extensive applications in robotics (providing protocols, e.g., for efficient patrolling and guarding of terrains, traversing terrains using groups of robots, etc.). We have already published several papers in this area (SIROCCO 2010 [79], SWAT 2010 [80], ESA 2011 [78]), building up the theoretical fundamentals of a new field, and already attracting the attention of a wider community of researchers working in robotics and AI.

Perspectives: Our goal is to explore applications of mobile agent techniques in domains of growing importance, namely, social networks and robotics. We are currently discussing applications of our techniques in problems of brand recognition on the web with a local industrial partner (Systonic KeepAlert), and other companies (through our research collaborators in Liverpool). We intend to undertake collaboration with European/American research labs and industrial partners.

5. Software

5.1. SimGrid

Participants: Przemyslaw Uznanski, Lionel Eyraud-Dubois [correspondant].

SimGrid (http://simgrid.gforge.inria.fr/) SimGrid is a toolkit that provides core functionalities for the simulation of distributed applications in heterogeneous distributed environments. The specific goal of the project is to facilitate research in the area of parallel and distributed large scale systems, such as Grids, P2P systems and clouds. Its use cases encompass heuristic evaluation, application prototyping or even real application development and tuning. It is based on experimentally validated models, and features very high scalability, which allows to perform very large scale simulations. It is used by over a hundred academic users all over the world, and has been used in about one hundred scientific articles.

CEPAGE has contributed to this software by participating in the management of the project and in many design decisions. We also implemented simpler models, based on our works in this area, allowing a better scalability while keeping a reasonable precision.

Software assessment : A-4, SO-4, SM-4, EM-3, SDL-5.
Contribution : DA-2, CD-2, MS-2, TPM-3.

5.2. Hubble

Participants: Ludovic Courtès, Nicolas Bonichon [correspondant].

Hubble is implemented in Scheme, using GNU Guile version 2. Details of the simulation, such as keeping track of processor occupation and network usage, are taken care of by SimGrid, a toolkit for the simulation of distributed applications in heterogeneous distributed environments.

The input to Hubble is an XML description of the DAG of build tasks. For each task, a build duration and the size in bytes of the build output are specified. For our evaluation purposes, we collected this data on a production system, the http://hydra.nixos.org/ build farm hosted at the Technical University of Delft. The DAG itself is the snapshot of the Nix Package Collection (Nixpkgs) corresponding to this data. Hubble has its own in-memory representation of the DAG in the form of a purely functional data structure.

The Nixpkgs DAG contains fixed-output nodes, i.e., nodes whose output is known in advance and does not require any computation. These nodes are typically downloads of source code from external web sites. The raw data collected on http://hydra.nixos.org/ specifies a non-zero duration for these nodes, which represents the time it took to perform the download. This duration info is irrelevant in our context, since they don't require any computation, and Hubble views these nodes as instantaneous.

See also the web page http://hubble.gforge.inria.fr/.

Software assessment: A-3, SO-3, SM-2, EM-1, SDL-2.
Contribution: DA-4, CD-4, MS-4, TPM-4.

5.3. Gengraph

Participant: Cyril Gavoille [correspondant].

This is a command-line tool for generating graphs. There are several output formats, includes the dot format from GraphViz. It generates also .pdf files for visualization. Several graph algorithms have been implemented (diameter, connectivity, treewidth, etc.) which can be tested on the graphs. The software has been originally designed for teaching purpose so that students can test their project algorithms on many non trivial families like random geometric graphs, graphs of given density, given treewidth. It is also used for research purpose, in particular the exhaustive search results in the Emilie Diot's thesis are based on `gengraph`. The program can filter a list of graphs based to many criteria, as for instance it can extract all graphs of a given list that are 2-connected, of diameter at least four, and that exclude some minor (or some induced subgraph).

Currently, more than 100 parametrized graph families are implemented, supporting simple operators like complementation, random edge/vertex removal, and others. The source has more than 10,000 lines including a command-line documentation of 2,000 lines. The single source file is available at http://dept-info.labri.fr/~gavoille/gengraph.c

Software assessment: A-3, SO-3, SM-2, EM-2, SDL-2.

Contribution: DA-4, CD-4, MS-4, TPM-4.

5.4. Bedibe

Participants: Lionel Eyraud-Dubois [correspondant], Przemyslaw Uznanski.

Bedibe (Benchmarking Distributed Bandwidth Estimation) is a software to compare different models for bandwidth estimation on the Internet, and their associated instantiation algorithms. The goal is to ease the development of new models and algorithms, and the comparison with existing solutions. The development of this software is just starting.

See also the web page http://bedibe.gforge.inria.fr/.

Software assessment : A-1-up2, SO-3, SM-1-up2, EM-2, SDL-1-up2.

5.5. MineWithRounds

Participants: Sofian Maabout [correspondant], Nicolas Hanusse.

The software implements a parallel algorithm aiming at computing *Borders* that's sets of maximal/minimal subsets of objects satisfying some anti-monotone condition. It is implemented in `C++` together with the `openMP` library to exploit multi-core machines. In its current status, it outperforms state of the art implementations addressing the Maximal Frequent Itemsets problem.

Software assessment: A-2, SO-4, SM-2, EM-2, SDL-2.
Contribution: DA-4, CD-4, MS-4, TPM-4.

6. New Results

6.1. Resource allocation and Scheduling

6.1.1. Divisible Load Scheduling

Participants: Olivier Beaumont, Nicolas Bonichon, Lionel Eyraud-Dubois.

Malleable tasks are jobs that can be scheduled with preemptions on a varying number of resources. In [22], we focus on the special case of work-preserving malleable tasks, for which the area of the allocated resources does not depend on the allocation and is equal to the sequential processing time. Moreover, we assume that the number of resources allocated to each task at each time instant is bounded. We consider both the clairvoyant and non-clairvoyant cases, and we focus on minimizing the weighted sum of completion times. In the weighted non-clairvoyant case, we propose an approximation algorithm whose ratio (2) is the same as in the unweighted non-clairvoyant case. In the clairvoyant case, we provide a normal form for the schedule of such malleable tasks, and prove that any valid schedule can be turned into this normal form, based only on the completion times of the tasks. We show that in these normal form schedules, the number of preemptions per task is bounded by 3 on average. At last, we analyze the performance of greedy schedules, and prove that optimal schedules are greedy for a special case of homogeneous instances. We conjecture that there exists an optimal greedy schedule for all instances, which would greatly simplify the study of this problem. Finally, we explore the complexity of the problem restricted to homogeneous instances, which is still open despite its very simple expression. (Joint work with Loris Marchal from ENS Lyon)

6.1.2. Scheduling for Distributed Continuous Integration

Participants: Olivier Beaumont, Nicolas Bonichon, Ludovic Courtès.

In [21], we consider the problem of schedul- ing a special kind of mixed data-parallel applications arising in the context of continuous integration. Continuous integration (CI) is a software engineering technique, which consists in re- building and testing interdependent software components as soon as developers modify them. The CI tool is able to provide quick feedback to the developers, which allows them to fix the bug soon after it has been introduced. The CI process can be described as a DAG where nodes represent package build tasks, and edges represent dependencies among these packages; build tasks themselves can in turn be run in parallel. Thus, CI can be viewed as a mixed data-parallel application. A crucial point for a successful CI process is its ability to provide quick feedback. Thus, makespan minimization is the main goal. Our contribution is twofold. First, we provide and analyze a large dataset corresponding to a build DAG. Second, we compare the performance of several scheduling heuristics on this dataset.

6.1.3. Resource Allocation in Clouds

Participants: Olivier Beaumont, Lionel Eyraud-Dubois, Hejer Rejeb.

In [14], we consider the problem of assigning a set of clients with demands to a set of servers with capacities and degree constraints. The goal is to find an allocation such that the number of clients assigned to a server is smaller than the server's degree and their overall demand is smaller than the server's capacity, while maximizing the overall throughput. This problem has several natural applications in the context of independent tasks scheduling or virtual machines allocation. We consider both the *offline* (when clients are known beforehand) and the *online* (when clients can join and leave the system at any time) versions of the problem. We first show that the degree constraint on the maximal number of clients that a server can handle is realistic in many contexts. Then, our main contribution is to prove that even if it makes the allocation problem more difficult (NP-Complete), a very small additive resource augmentation on the servers degree is enough to find in polynomial time a solution that achieves at least the optimal throughput. After a set of theoretical results on the complexity of the offline and online versions of the problem, we propose several other greedy heuristics to solve the online problem and we compare the *performance* (in terms of throughput) and the *cost* (in terms of disconnections and reconnections) of all proposed algorithms through a set of extensive simulation results. (Joint work with Christopher Thraves-Caros, University of Madrid)

6.1.4. Non Linear Divisible Load Scheduling

Participants: Olivier Beaumont, Hubert Larchevêque.

Divisible Load Theory (DLT) has received a lot of attention in the past decade. A divisible load is a perfect parallel task, that can be split arbitrarily and executed in parallel on a set of possibly heterogeneous resources. The success of DLT is strongly related to the existence of many optimal resource allocation and scheduling algorithms, what strongly differs from general scheduling theory. Moreover, recently, close relationships have been underlined between DLT, that provides a fruitful theoretical framework for scheduling jobs on heterogeneous platforms, and MapReduce, that provides a simple and efficient programming framework to deploy applications on large scale distributed platforms. The success of both have suggested to extend their framework to non-linear complexity tasks. In [24], we show that both DLT and MapReduce are better suited to workloads with linear complexity. In particular, we prove that divisible load theory cannot directly be applied to quadratic workloads, such as it has been proposed recently. We precisely state the limits for classical DLT studies and we review and propose solutions based on a careful preparation of the dataset and clever data partitioning algorithms. In particular, through simulations, we show the possible impact of this approach on the volume of communications generated by MapReduce, in the context of Matrix Multiplication and Outer Product algorithms. (Joint work with Loris Marchal from ENS Lyon)

6.1.5. Reliable Service Allocation in Clouds

Participants: Olivier Beaumont, Lionel Eyraud-Dubois, Hubert Larchevêque.

In [23], we consider several reliability problems that arise when allocating applications to processing resources in a Cloud computing platform. More specifically, we assume on the one hand that each computing resource is associated to a capacity constraint and to a probability of failure. On the other hand, we assume that each service runs as a set of independent instances of identical Virtual Machines, and that the Service Level Agreement between the Cloud provider and the client states that a minimal number of instances of the service should run with a given probability. In this context, given the capacity and failure probabilities of the machines, and the capacity and reliability demands of the services, the question for the cloud provider is to find an allocation of the instances of the services (possibly using replication) onto machines satisfying all types of constraints during a given time period. The goal of this work is to assess the impact of the reliability constraint on the complexity of resource allocation problems. We consider several variants of this problem, depending on the number of services and whether their reliability demand is individual or global. We prove several fundamental complexity results ($\#$P' and NP-completeness results) and we provide several optimal and approximation algorithms. In particular, we prove that a basic randomized allocation algorithm, that is easy to implement, provides optimal or quasi-optimal results in several contexts, and we show through simulations that it also achieves very good results in more general settings.

6.1.6. Optimizing Resource allocation while handling SLA violations in Cloud Computing platforms

Participants: Lionel Eyraud-Dubois, Hubert Larchevêque.

In [29], we study a resource allocation problem in the context of Cloud Computing, where a set of Virtual Machines (VM) has to be placed on a set of Physical Machines (PM). Each VM has a given demand (e.g. CPU demand), and each PM has a capacity. However, each VM only uses a fraction of its demand. The aim is to exploit the difference between the demand of the VM and its real utilization of the resources, to exploit the capacities of the PMs as much as possible. Moreover, the real consumption of the VMs can change over time (while staying under its original demand), implying sometimes expensive "SLA violations", corresponding to some VM's consumption not satisfied because of overloaded PMs. Thus, while optimizing the global resource utilization of the PMs, it is necessary to ensure that at any moment a VM's need evolves, a few number of migrations (moving a VM from PM to PM) is sufficient to find a new configuration in which all the VMs' consumptions are satisfied. We modelize this problem using a fully dynamic bin packing approach and we present an algorithm ensuring a global utilization of the resources of 66%. Moreover, each time a PM is overloaded at most one migration is necessary to fall back in a configuration with no overloaded PM, and only 3 different PMs are concerned by required migrations that may occur to keep the global resource utilization correct. This allows the platform to be highly resilient to a great number of changes.

6.2. Compact Routing

6.2.1. Compact routing with forbidden-set in planar graphs

Participant: Cyril Gavoille.

In [20], we consider fully dynamic $(1+\varepsilon)$ distance oracles and $(1+\varepsilon)$ forbidden-set labeling schemes for planar graphs. For a given n-vertex planar graph G with edge weights drawn from $[1, M]$ and parameter $\varepsilon > 0$, our forbidden-set labeling scheme uses labels of length $\lambda = O(\varepsilon^{-1}\log^2 n \log{(nM)} \cdot \max{logn})$. Given the labels of two vertices s and t and of a set F of faulty vertices/edges, our scheme approximates the distance between s and t in $G \setminus F$ with stretch $(1+\varepsilon)$, in $O(|F|^2\lambda)$ time.

We then present a general method to transform $(1+\varepsilon)$ forbidden-set labeling schemas into a fully dynamic $(1+\varepsilon)$ distance oracle. Our fully dynamic $(1+\varepsilon)$ distance oracle is of size $O(n\log n \cdot \max{logn})$ and has $\widetilde{O}(n^{1/2})$ query and update time, both the query and the update time are worst case. This improves on the best previously known $(1+\varepsilon)$ dynamic distance oracle for planar graphs, which has worst case query time $\widetilde{O}(n^{2/3})$ and amortized update time of $\widetilde{O}(n^{2/3})$.

Our $(1+\varepsilon)$ forbidden-set labeling scheme can also be extended into a forbidden-set labeled routing scheme with stretch $(1+\varepsilon)$.

6.2.2. Planar Spanner of geometric graphs

Participants: Nicolas Bonichon, Cyril Gavoille, Nicolas Hanusse.

In [26], we determine the stretch factor of L_1-Delaunay and L_∞-Delaunay triangulations, and we show that this stretch is $\sqrt{4+2\sqrt{2}} \approx 2.61$. Between any two points x, y of such triangulations, we construct a path whose length is no more than $\sqrt{4+2\sqrt{2}}$ times the Euclidean distance between x and y, and this bound is best possible. This definitively improves the 25-year old bound of $\sqrt{10}$ by Chew (SoCG '86).

To the best of our knowledge, this is the first time the stretch factor of the well-studied L_p-Delaunay triangulations, for any real $p \geq 1$, is determined exactly.

6.3. Mobile Agents

6.3.1. More efficient periodic traversal in anonymous undirected graphs

Participants: David Ilcinkas, Ralf Klasing.

In [15], we consider the problem of *periodic graph exploration* in which a mobile entity with constant memory, *an agent*, has to visit all n nodes of an input simple, connected, undirected graph in a periodic manner. Graphs are assumed to be anonymous, that is, nodes are unlabeled. While visiting a node, the agent may distinguish between the edges incident to it; for each node v, the endpoints of the edges incident to v are uniquely identified by different integer labels called *port numbers*. We are interested in algorithms for assigning the port numbers together with traversal algorithms for agents using these port numbers to obtain short traversal periods.

Periodic graph exploration is unsolvable if the port numbers are set arbitrarily; see Budach (1978). However, surprisingly small periods can be achieved by carefully assigning the port numbers. Dobrev *et al.* (2005) described an algorithm for assigning port numbers and an oblivious agent (i.e., an agent with no memory) using it, such that the agent explores any graph with n nodes within the period $10n$. When the agent has access to a constant number of memory bits, the optimal length of the period was proved in Gasieniec *et al.* (2008) to be no more than $3.75n - 2$ (using a different assignment of the port numbers and a different traversal algorithm). In our work, we improve both these bounds. More precisely, we show how to achieve a period length of at most $(4 + \frac{1}{3})n - 4$ for oblivious agents and a period length of at most $3.5n - 2$ for agents with constant memory. To obtain our results, we introduce a new, fast graph decomposition technique called a *three-layer partition* that may also be useful for solving other graph problems in the future. Finally, we present the first non-trivial lower bound, $2.8n - 2$, on the period length for the oblivious case.

6.3.2. Gathering of Robots on Anonymous Grids without Multiplicity Detection

Participant: Ralf Klasing.

In [28] , we study the gathering problem on grid networks. A team of robots placed at different nodes of a grid have to meet at some node and remain there. Robots operate in Look-Compute-Move cycles; in one cycle, a robot perceives the current configuration in terms of occupied nodes (Look), decides whether to move towards one of its neighbors (Compute), and in the positive case makes the computed move instantaneously (Move). Cycles are performed asynchronously for each robot. The problem has been deeply studied for the case of ring networks. However, the known techniques used on rings cannot be directly extended to grids. Moreover, on rings, another assumption concerning the so-called *multiplicity detection* capability was required in order to accomplish the gathering task. That is, a robot is able to detect during its Look operation whether a node is empty, or occupied by one robot, or occupied by an undefined number of robots greater than one.

In our work, we provide a full characterization about gatherable configurations for grids. In particular, we show that in this case, the multiplicity detection is not required. Very interestingly, sometimes the problem appears trivial, as it is for the case of grids with both odd sides, while sometimes the involved techniques require new insights with respect to the well-studied ring case. Moreover, our results reveal the importance of a structure like the grid that allows to overcome the multiplicity detection with respect to the ring case.

7. Partnerships and Cooperations

7.1. Regional Initiatives

+ **CRA Region** (participants: CEPAGE). This project, entitled "Services for large-scale distributed platforms", is an effort for the designing efficient algorithms for clustering and discovering resources in large scale distributed networks. This project provided the funding for the PhD thesis of Hubert Larcheveque.
+ **CRA Region** (participants: CEPAGE, RUNTIME (Bordeaux)). This project, entitled "Performance modeling for heterogeneous platforms", is an effort for the modeling of the behavior of applications on two different types of platforms: multicore architectures within the RUNTIME team, and large scale platforms within CEPAGE. This project provides the funding for the PhD thesis of Przemyslaw Uznanski.

7.2. National Initiatives

- **ANR ALADDIN** (Algorithm Design and Analysis for Implicitly and Incompletely Defined Interaction Networks; GANG and CEPAGE project-teams): the members of Cepage have been participating to the ANR project "blanc" (i.e. fundamental research) about the fundamental aspects of large interaction networks enabling massive distributed storage, efficient decentralized information retrieval, quick inter-user exchanges, and/or rapid information dissemination. The project is mostly oriented towards the design and analysis of algorithms for these (logical) networks, by taking into account proper ties inherent to the underlying infrastructures upon which they are built. The infrastructures and/or overlays considered in this project are selected from different contexts, including communication networks (from Internet to sensor networks), and societal networks (from the Web to P2P networks).
- **ANR USS-SIMGRID** (Ultra Scalable Simulations with SimGrid; participants: AlGorille (LORIA, Nancy), ASAP (Saclay), CEPAGE, Univ. of Hawai'i, GRAAL (LIP, ENS Lyon), MESCAL (Grenoble), MASCOTTE (Sophia Antipolis)). The members of CEPAGE were part of this project (2008-2011), whose goal was to extend the SimGrid simulation framework, originally developed for HPC, to provide a reasonable and quantifiable level of accuracy for the simulation of large scale application. This allowed to attend both the rising need for scalability of the HPC community and the need for simulation accuracy of the distributed computing community. SimGrid was extended to provide a family of models which offer different levels of accuracy at different simulation scales.
- **ANR SONGS** (Simulation of Next Generation Systems; participants: AlGorille (LORIA, Nancy), MESCAL (Grenoble), GRAAL (ENS Lyon), IN2P3 (Lyon), CEPAGE, HiePACS, RUNTIME (Bordeaux), LSIIT (Strasbourg), ASCOLA (Nantes), MASCOTTE, MODALIS (Sophia Antipolis)). This project started in 2012 as a follow-up of the USS-SIMGRID project. The aim is to further extend the domain of SimGrid, by designing a unified simulation framework for the four application domains: Grids, Peer-to-Peer systems, High Performance Computing, and Cloud systems. Achieving this goal mandates careful representation and modeling of the underlying concepts presented by each domain (memory, disks, energy, network and volatility) and of the interfaces specific to each domain. It also requires a transversal work on the simulation framework itself. CEPAGE is actively involved in this project, both for the peer-to-peer use cases and for the coordination of the modeling effort of the project.
- **ANR Displexity** (Calcul DIStribué: calculabilité et comPLEXITé; participants: CEPAGE, GANG and ASAP projects). The main goal of DISPLEXITY is to establish the scientific foundations of a theory of calculability and complexity for distributed computing. Displexity started in 2012.
- **ANR IDEA** ANR program “defis”: project IDEA (2009-2012). The goal of this ANR is the study of identifying codes in evolving graphs. Ralf Klasing is the overall leader of the project.

- **ANR “Jeunes chercheurs” EGOS - Embedded Graphs and their Oriented Structures** (2012-2014) (see http://www.lirmm.fr/egos/)

  Participants: CEPAGE/LaBRI(Bordeaux) LIRMM(Montpellier), LIX(Palaiseau) The goal of this project is the study oriented structures on graphs of arbitrary genus.
- **AMADEUS** (CNRS funding on “BIG DATA”: 2012-): Analysis of MAssive Data in Earth and Universe Sciences. This a multidisciplinary research project between computer science teams (LIRMM: University of Montpellier, LIF: University of Marseille) and CEPAGE), earth and climate science (CEREGE: Montpellier and IRD: Aix) and astronomy (LAM: University of Marseille). The aim of the project is to propose effective techniques for mining large data by essentially using distributed computing, visualization, summarization and approximation.

7.3. European Initiatives

7.3.1. EULER

Title: EULER (Experimental UpdateLess Evolutive Routing)

Type: COOPERATION (ICT)

Defi: Future Internet Experimental Facility and Experimentally-driven Research

Instrument: Specific Targeted Research Project (STREP)

Duration: October 2010 - September 2013

Coordinator: ALCATEL-LUCENT (Belgium)

Others partners:
Alcatel-Lucent Bell, Antwerpen, Belgium
3 projects from Inria: CEPAGE, GANG and MASCOTTE, France
Interdisciplinary Institute for Broadband Technology (IBBT),Belgium
Laboratoire d’Informatique de Paris 6 (LIP6), Université Pierre Marie Curie (UPMC), France
Department of Mathematical Engineering (INMA) Université Catholique de Louvain, Belgium
RACTI, Research Academic Computer Technology Institute University of Patras, Greece
CAT, Catalan Consortium: Universitat Politecnica de Catalunya, Barcelona and University of Girona, Spain

See also: http://www-sop.inria.fr/mascotte/EULER/wiki/

Abstract: The title of this study is "Dynamic Compact Routing Scheme". The aim of this projet is to develop new routing schemes achieving better performances than current BGP protocols. The problems faced by the inter-domain routing protocol of the Internet are numerous:

The underlying network is dynamic: many observations of bad configurations show the instability of BGP;

BGP does not scale well: the convergence time toward a legal configuration is too long, the size of routing tables is proportional to the number of nodes of network (the network size is multiplied by 1.25 each year);

The impact of the policies is so important that the many packets can oscillated between two Autonomous Systems.

In this collaboration, we mainly focus on the scalability properties that a new routing protocol should guarantee. The main measures are the size of the local routing tables, and the time (or message complexity) to update or to generate such tables. The design of schemes achieving sub-linear space per routers, say in n where n is the number of AS routers, is the main challenge. The target networks are AS-network like with more than 100,000 nodes. This projet, in collaboration with the MASCOTE Inria-project in Nice Sophia-Antipolis, makes the use of simulation, developed at both sites.

7.3.2. Collaborations in European Programs, except FP7

Program: European COST

Project acronym: Complex HPC IC0805.

Project title: Open Network for High-Performance Computing on Complex Environments

Duration: 2010-2013

Coordinator: Inria

Other partners: 26 countries, see list at http://www.cost.eu/domains_actions/ict/Actions/IC0805?parties

Abstract: The main objective of this COST action is to coordinate European groups working on the use of heterogeneous and hierarchical systems for HPC as well as the development of collaborative activities among the involved research groups (http://complexhpc.org/index.php).

7.4. International Initiatives

7.4.1. Participation In International Programs

- **Royal Society Grant with the University of Liverpool.** International Joint Project, 2011-2013, entitled “SEarch, RENdezvous and Explore (SERENE)”, on foundations of mobile agent computing, in collaboration with the Department of Computer Science, University of Liverpool. Funded by the Royal Society, U.K. Principal investigator on the UK side: Leszek Gasieniec. Ralf Klasing is the principal investigator on the French side.

  Participants: Nicolas Hanusse, David Ilcinkas, Ralf Klasing, Adrian Kosowski.
- **Spanish program CLOUDS**: Cloud Computing for Scalable, Reliable and Ubiquitous Services (http://lsd.ls.fi.upm.es/clouds). This is a large scale program which aims at advancing research in the area of Cloud Computing. CEPAGE is more particularly in contact with the LaDyr team of Univ. Rey Juan Carlos in Madrid, on the topic of resource allocation problems for Cloud providers.

  Participants: Olivier Beaumont, Lionel Eyraud-Dubois.
- **Collaboration with Canada.**

  Members of CEPAGE have a long-standing collaboration with researchers from the Chair of Distributed Computing at the University of Quebec in Outaouais and the Department of Computer Science at Carleton University. Sources of financing include: personal NSERC grants of Canadian professors (Prof. Andrzej Pelc, Prof. Jurek Czyzowicz, Prof. Evangelos Kranakis), funding from other Canadian grant agencies (a travel grant from Mitacs Inc.), and University of Bordeaux funding (a 3-month invited professorship for Prof. Jurek Czyzowicz).

  Participants: David Ilcinkas, Ralf Klasing, Adrian Kosowski.
- **Collaboration with Chile.**

  Adrian Kosowski is a foreign partner of the Chilean ministry grant (ANILLO CONICYT programme) entitled “Mathematical modeling for industrial and management science applications: a multidisciplinary approach”. The Project Director is Eric Goles from Universidad Adolfo Ibañez, and collaborating researchers on the Chilean side include Karol Suchan and Ivan Rappaport. The collaboration has led to 2 joint papers.

  Participants: Adrian Kosowski.

7.5. International Research Visitors

7.5.1. Visits of International Scientists

7.5.1.1. Visits to Cepage Members

- Ljubomir Perkovic, De Paul University Chicago, (September 2011 – June 2012)

- Prosenjit Bose, Carleton University Ottawa, (25/11/12 – 29/11/12)
- George Mertzios, Durham University, UK, (15/06/12 – 14/07/12)
- Leszek Gasieniec, University of Liverpool, UK, (08/06 – 22/06/12)
- Jurek Czyzowicz, Université du Québec, Canada, (08/06 – 22/06/12)
- Darek Dareniowski, Gdansk University of Technology, Poland, (08/06 – 28/06/2012)
- Miroslaw Korzeniowski, Technical University of Wroclaw, (March 2012 – September 2012)

7.5.1.2. Visits of Cepage Members

- Cyril Gavoille, MicroSoft Research, Mountainview, CA, two weeks in April 2012.

8. Dissemination

8.1. Scientific Animation

8.1.1. Editorial Work

- Ralf Klasing is Associate Editor for
  - Algorithmic Operations Research (since May 2007),
  - Parallel Processing Letters (since August 2007),
  - Networks (since September 2007),
  - Computing and Informatics (since January 2008),
  - **Theoretical Computer Science (since December 2009),**
  - Fundamenta Informaticae (since January 2010),
  - Discrete Applied Mathematics (since February 2010),
  - Wireless Networks (since May 2010),
  - Journal of Interconnection Networks (since November 2010).
- Olivier Beaumont is Associate Editor for
  - **IEEE Transactions on Parallel and Distributed Systems (since June 2010).**

8.1.2. Steering Committees

- Ralf Klasing is a member of the Steering Commitee of the *International Colloquium on Structural Information and Communication Complexity (SIROCCO).*

8.1.3. Organizing Commitees

- Ralf Klasing was a member of the organization committee of the *Bordeaux Graph Workshop*, Bordeaux, France, November 21-24, 2012.

8.1.4. Program Committees – Chair

- Nicolas Hanusse was the Conference Chair of *14th French Conference of Communications in Network (Algotel 2012)*
- Nicolas Hanusse was the Conference Chair of *International Conference on Algorithm (FUN 2012)*
- Ralf Klasing was the Conference Chair of the *11th International Symposium on Experimental Algorithms (SEA 2012)*, Bordeaux, France, June 7-9, 2012.
- Philippe Duchon was the Conference Chair of the *The 8th edition of the conference GASCom on random generation of combinatorial structures*, Bordeaux, France, June 7-9, 2012.
- Sofian Maabout was the conference Chair of *8th French Conference on Data Warehouses and OLAP (EDA2012)*, Bordeaux, June 12-13, 2012.

- Olivier Beaumont will be the Conference Chair (Algorithms Track) of the *42th International Conference on Parallel Processing*, ICPP 2013, Lyon, France, 2013

8.1.5. Program Committees

- Olivier Beaumont
  - ICPP 2012, 41st International Conference on Parallel Processing, Pittsburgh, PA, USA, September 10-13, 2012
  - IPDPS 2012, 26th IEEE International Parallel & Distributed Processing Symposium, May 21-25, 2012, Shanghai, China
  - HCW 2012, 20th International Heterogeneity in Computing Workshop, May 21-25, 2012, Shanghai, China
  - ISCIS 2012, 27th International Symposium on Computer and Information Sciences, Paris, France
  - CSE 2012, 15th IEE International Conference on Computational Science and Engineering, Paphos, Cyprus
- Nicolas Bonichon
  - AlgoTel 2012 (May 29- June 1, La Grande Motte, France) Rencontres Francophones sur les aspects Algorithmiques des Télécommunications
- Cyril Gavoille
  - DISC 2012 (Oct. 16-18, Salvador, Brazil), International Symposium on Distributed Computing
  - ISAAC 2012 (Dec. 19-21, Taipei, Taiwan), International Symposium on Algorithms and Computation
  - ICDCN 2012 (Jan. 3-6, Hong-Kong), International Conference on Distributed Computing and Networking
- David Ilcinkas
  - FOMC 2012 (8^{th} ACM SIGACT/SIGMOBILE International Workshop on Foundations of Mobile Computing)
  - **OPODIS 2012 (16^{th} International Conference On Principles Of DIstributed Systems)**
- Ralf Klasing
  - SEA 2012, 11th International Symposium on Experimental Algorithms, Bordeaux, France, June 7-9, 2012
  - ADHOC-NOW 2012, 11th International Conference on Ad Hoc Networks and Wireless, July 9-11, 2012, Belgrade, Serbia
  - FOMC 2012, 8th ACM SIGACT/SIGMOBILE International Workshop on Foundations of Mobile Computing (formerly known as DIALM-POMC), July 19, 2012, Madeira, Portugal
  - IWOCA 2012, 23rd International Workshop on Combinatorial Algorithms, July 19-21, 2012, Kalasalingam University, Anand Nagar, Krishnankoil, Tamil Nadu, India
- Lionel Eyraud-Dubois
  - Cluster 2012, *IEEE Conference on Cluster Comuting*, September 24-28, Beijing, China.
- Adrian Kosowski
  - SEA 2012, 11th International Symposium on Experimental Algorithms, Bordeaux, France, June 7-9, 2012.
  - SSS 2012, 14th International Symposium on Stabilization, Safety, and Security of Distributed Systems, Toronto, September 24-27, 2012 .

  - ICDCN 2012, 13th International Conference on Distributed Computing and Networking, January 3-6, 2012, Hong Kong, China.
  - MFCS 2012, 37th International Symposium on Mathematical Foundations of Computer Science, August 27-31, 2012, Bratislava, Slovakia.
- Sofian Maabout
  - ICWIT 2012, 4th Int. Conf. on Web and Information Technologies, Sidi Bel Abbes, Algeria, April 29-30, 2012.

8.1.6. Research Administration

8.1.6.1. Main Administrative Duties

- Olivier Beaumont is the Scientific Deputy of Inria Bordeaux Sud-Ouest (since September 2011)
- Nicolas Hanusse is the Head of the Doctoral School (Computer Science) of the University of Bordeaux (since 2010)
- Cyril Gavoille was a Deputy Director of LaBRI laboratory (2008–2011)
- Ralf Klasing is responsible of the "Combinatorics and Algorithms" Team of the LaBRI. (since November 2010)
- Nicolas Hanusse is responsible of the "Distributed Algorithms" Group in LaBRI. (since November 2010)
- Philippe Duchon is the Head of the University Bordeaux 1 Computer Science Master program (since 2010)
- Ralf Klasing is responsible for the International Relations of the LaBRI. (since January 2009)
- Sofian Maabout is responsible of the second year of the "Information Systems for Health Care" Master (since 2010).
- Ralf Klasing is a member of the Evaluation Committee of the programme "ANR DEFIS".
- David Ilcinkas is a member of the *Conseil de Laboratoire* of the LaBRI.

8.2. Teaching - Supervision - Juries

8.2.1. Teaching

Master : Communication and Routing (last year of engineering school ENSEIRB, 2012) O. Beaumont, N. Bonichon, L. Eyraud, N. Hanusse, R. Klasing, A. Kosowski (16h)

Master : Communication Algorithms in Networks (2nd year MASTER "Algorithms and Formal Methods", University of Bordeaux, 2012) R. Klasing (24h)

Master : Search Engines (2nd year of engineering school ENSEIRB, 2012) O. Beaumont (20h)

Master : Distributed Computing (2nd year MASTER "Réseaux, Systèmes et Mobilité"), (24h)

8.2.2. Supervision

PhD :

- Florent Foucaud, Aspects combinatoires et algorithmiques des codes identifiants dans les graphes, Ralf Klasing, André Raspaud, Université Bordeaux 1, 10.12.2012
- Godfroy Quentin, From spanners to multipath spanners, Cyril Gavoille, Université Bordeaux 1, dec. 2012.

PhD in progress :

- Przemyslaw Uznanski, Communication modeling on large scale platforms, November 2010, Olivier Beaumont, Nicolas Bonichon, Lionel Eyraud-Dubois, Université Bordeaux 1.

- Ahmed Wade, Mobile agent protocols for dynamic networks, February 2011, David Ilcinkas, Ralf Klasing, Université Bordeaux 1.
- Dominik Pajak, Multi-agent protocols for efficient graph exploration, October 2011, Adrian Kosowski, Ralf Klasing, Université Bordeaux 1.
- Pierre Halftermeyer, Structuration des graphes et étiquetages compacts, October 2010, Cyril Gavoille, Université Bordeaux 1.

8.2.3. Juries

- Ralf Klasing was examinateur and president of the PhD committee of Thomas Morsellino (University Bordeaux 1, 25.9.2012)
- Ralf Klasing was external PhD reviewer and member of the PhD committee in the Ph.D. defense of Aline Parreau (University of Grenoble, 5.7.2012)
- Ralf Klasing was external PhD reviewer and member of the PhD committee in the Ph.D. defense of Marwane Bouznif (University of Grenoble, 4.7.2012)
- Cyril Gavoille was external reviewer for the PhD thesis of Mikaila Toko Worou, Université Nice Sophia-Antipolis, Outils algorithmiques pour la détection des communautés dans les réseaux (12/2012).
- Cyril Gavoille was was external reviewer for the PhD thesis of Jean-François Couturier, Université Paul Verlaine (Metz), LITA, Algorithmes exacts et exponentiels sur les graphes: énumération, comptage et optimisation (12/2012).
- Cyril Gavoille was (anonymous) external reviewer for the PhD thesis of Université des Sciences et de la Technologie Houari Boumédiène in Algeria (11/2012).
- Cyril Gavoille was president of the committee and examinator for the PhD thesis of Nicolas Delfosse, Université de Bordeaux, IMB, Constructions et performances de codes LDPC quantiques (12/2012).
- Nicolas Hanusse was examinator for the PhD thesis of Dominique Dion, Dyname d'évolution de graphes de cooccurences lexicales: application à l'analyse de comptes-rendus en prévention spécialisée entre 1972 et 2010.
- Nicolas Hanusse was president of the committee and examinator for the PhD thesis of Quention Godfroy, Multipath Spanners, Université de Bordeaux (11/2012)
- Nicolas Hanusse was external reviewer and examinator for the "Habilitation of Research Direction" of Jean-Loup Guillaume, Déterminisme et non-déterminisme au service de la détection de communautés dynamiques, Paris 6 University (11/2012)
- Nicolas Hanusse was president of the committee and examinator for the PhD thesis of Asma Ben Zakour, Extraction des utilisations typiques à partir de données hétérogènes historisées en vue d'optimiser la maintenance d'une flotte de véhicules , Université de Bordeaux (06/2012)

9. Bibliography

Major publications by the team in recent years

[1] O. BEAUMONT, H. REJEB. *On the Importance of Bandwidth Control Mechanisms for Scheduling on Large Scale Heterogeneous Platforms*, in "24th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2010)", Atlanta, United States, April 2010, http://hal.inria.fr/inria-00444585/en.

[2] N. BONICHON, C. GAVOILLE, N. HANUSSE, L. PERKOVIC. *Plane Spanners of Maximum Degree Six*, in "Automata, Languages and Programming 37th International Colloquium, ICALP 2010, Bordeaux, France, July 6-10, 2010, Proceedings, Part I", France, 2010, vol. 6198, p. 19-30 [*DOI :* 10.1007/978-3-642-14165-2_3], http://hal.archives-ouvertes.fr/hal-00534212/en/.

[3] C. COOPER, D. ILCINKAS, R. KLASING, A. KOSOWSKI. *Derandomizing random walks in undirected graphs using locally fair exploration strategies*, in "Distributed Computing", 2011, vol. 24, n^o 2, p. 91-99, See paper for details. [*DOI :* 10.1007/s00446-011-0138-4], http://hal.archives-ouvertes.fr/hal-00638229/en/.

[4] J. CZYZOWICZ, A. KOSOWSKI, A. PELC. *How to meet when you forget: log-space rendezvous in arbitrary graphs*, in "29th Annual ACM Symposium on Principles of Distributed Computing (PODC)", ACM Press, July 2010, p. 450-459.

[5] P. DUCHON, N. HANUSSE, E. LEBHAR. *Towards small world emergence*, in "Proceedings of the eighteenth annual ACM symposium on Parallelism in algorithms and architectures", ACM Press New York, NY, USA, 2006, p. 225–232.

[6] P. DUCHON, N. HANUSSE, N. SAHEB, A. ZEMMARI. *Broadcast in the rendezvous model*, in "Information and Computation", 2006, vol. 204, n^o 5, p. 697–712.

[7] L. EYRAUD-DUBOIS, G. MOUNIÉ, D. TRYSTRAM. *Analysis of Scheduling Algorithms with Reservations*, in "IEEE International Parallel and Distributed Processing Symposium, IPDPS 2007", Long Beach, California, USA, IEEE, 2007, p. 1-8, http://dx.doi.org/10.1109/IPDPS.2007.370304.

[8] P. FLOCCHINI, D. ILCINKAS, N. SANTORO. *Ping Pong in Dangerous Graphs: Optimal Black Hole Search with Pure Tokens*, in "Proceedings of the 22nd International Symposium on Distributed Computing", France, G. TAUBENFELD (editor), Lecture Notes in Computer Science, Springer Berlin / Heidelberg, September 2008, vol. 5218, p. 227-241, This work was done during the stay of David Ilcinkas at the University of Ottawa, as a postdoctoral fellow. Paola Flocchini was partially supported by the University Research Chair of the University of Ottawa. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada under Discovery grants. [*DOI :* 10.1007/978-3-540-87779-0_16], http://hal.archives-ouvertes.fr/hal-00341523/en/.

[9] P. FRAIGNIAUD, C. GAVOILLE, D. ILCINKAS, A. PELC. *Distributed Computing with Advice: Information Sensitivity of Graph Coloring*, in "ICALP", 2007, p. 231-242.

[10] P. FRAIGNIAUD, C. GAVOILLE, A. KOSOWSKI, E. LEBHAR, Z. LOTKER. *Universal augmentation schemes for network navigability: overcoming the $\sqrt{n}$-barrier*, in "Proceedings of the nineteenth annual ACM symposium on parallelism and architectures", San Diego, California, États-Unis, ACM, 2007, p. 1-7, COST Action 295 "Dynamo", ACI project "Fragile", Inria project "Grand Large", http://hal.archives-ouvertes.fr/hal-00155186/en/.

[11] E. GARNAUD, S. MAABOUT, M. MOSBAH. *Using Functional Dependencies for Reducing the Size of a Data Cube*, in "Proceedings of the 7th Int. Symp. on Foundations of Information and Knowledge Systems (FoIKS)", LNCS, Springer, 2012, 20.

[12] L. GASIENIEC, R. KLASING, M. RUSSELL, A. NAVARRA, X. ZHANG. *Fast Periodic Graph Exploration with Constant Memory*, in "Journal of Computer and System Sciences", 2008, vol. 74, n^o 5, p. 808–822 [*DOI :* 10.1016/j.jcss.2007.09.004], http://hal.archives-ouvertes.fr/hal-00306995/en/.

[13] J. HROMKOVIČ, P. KANAREK, R. KLASING, K. LORYS, W. UNGER, H. WAGENER. *On the Size of Permutation Networks and Consequences for Efficient Simulation of Hypercube Algorithms on Bounded-Degree Networks*, in "SIAM Journal on Discrete Mathematics", 2009, vol. 23, n^o 3, p. 1612–1645, http://hal.archives-ouvertes.fr/hal-00402764/en/.

Publications of the year

Articles in International Peer-Reviewed Journals

[14] O. BEAUMONT, L. EYRAUD-DUBOIS, H. REJEB, C. THRAVES-CARO. *Heterogeneous Resource Allocation under Degree Constraints*, in "IEEE Transactions on Parallel and Distributed Systems", 2012 [*DOI :* 10.1109/TPDS.2012.175], http://hal.inria.fr/hal-00771773.

[15] J. CZYZOWICZ, S. DOBREV, L. GASIENIEC, D. ILCINKAS, J. JANSSON, R. KLASING, I. LIGNOS, R. MARTIN, K. SADAKANE, W.-K. SUNG. *More efficient periodic traversal in anonymous undirected graphs*, in "Theoretical Computer Science", July 2012, vol. 444, p. 60-76 [*DOI :* 10.1016/J.TCS.2012.01.035], http://hal.inria.fr/hal-00726767.

[16] J. CZYZOWICZ, A. KOSOWSKI, A. PELC. *How to meet when you forget: log-space rendezvous in arbitrary graphs*, in "Distributed Computing", 2012, vol. 25, n^o 2, p. 165-178 [*DOI :* 10.1007/S00446-011-0141-9], http://hal.inria.fr/hal-00725875.

[17] P. FLOCCHINI, D. ILCINKAS, N. SANTORO. *Ping Pong in Dangerous Graphs: Optimal Black Hole Search with Pebbles*, in "Algorithmica", April 2012, vol. 62, n^o 3-4, p. 1006-1033 [*DOI :* 10.1007/S00453-011-9496-3], http://hal.inria.fr/hal-00726790.

[18] F. FOUCAUD, R. KLASING, A. KOSOWSKI, A. RASPAUD. *On the size of identifying codes in triangle-free graphs*, in "Discrete Applied Mathematics", July 2012, vol. 160, n^o 10-11, p. 1532-1546 [*DOI :* 10.1016/J.DAM.2012.02.009], http://hal.inria.fr/hal-00530172.

[19] A. KOSOWSKI, A. NAVARRA. *Graph Decomposition for Memoryless Periodic Exploration*, in "Algorithmica", 2012, vol. 63, n^o 1-2, p. 26-38 [*DOI :* 10.1007/S00453-011-9518-1], http://hal.inria.fr/hal-00725922.

International Conferences with Proceedings

[20] I. ABRAHAM, S. CHECHIK, C. GAVOILLE. *Fully dynamic approximate distance oracles for planar graphs via forbidden-set distance labels*, in "44th Annual ACM Symposium on Theory of Computing (STOC)", New-York, United States, ACM (editor), May 2012, p. 1199-1217 [*DOI :* 10.1145/2213977.2214084], http://hal.inria.fr/hal-00725839.

[21] O. BEAUMONT, N. BONICHON, L. COURTÈS, X. HANIN, E. DOLSTRA. *Mixed Data-Parallel Scheduling for Distributed Continuous Integration*, in "Heterogeneity in Computing Workshop, in IPDPS 2012", Shangaï, China, IEEE (editor), 2012, http://hal.inria.fr/hal-00684220.

[22] O. BEAUMONT, N. BONICHON, L. EYRAUD-DUBOIS, L. MARCHAL. *Minimizing Weighted Mean Completion Time for Malleable Tasks Scheduling*, in "IPDPS 2012, 26th IEEE International Parallel & Distributed Processing Symposium", Shangai, China, IEEE (editor), 2012, http://hal.inria.fr/inria-00564056.

[23] O. BEAUMONT, L. EYRAUD-DUBOIS, H. LARCHEVÊQUE. *Reliable Service Allocation in Clouds*, in "PDPS 2013 - 27th IEEE International Parallel & Distributed Processing Symposium", Boston, États-Unis, 2013, http://hal.inria.fr/hal-00743524.

[24] O. BEAUMONT, H. LARCHEVÊQUE, L. MARCHAL. *Non Linear Divisible Loads: There is No Free Lunch*, in "IPDPS 2013, 27th IEEE International Parallel & Distributed Processing Symposium", Boston, États-Unis, IEEE, 2013, http://hal.inria.fr/hal-00771640.

[25] F. BECKER, A. KOSOWSKI, N. NISSE, I. RAPAPORT, K. SUCHAN. *Allowing Each Node to Communicate Only Once in a Distributed System: Shared Whiteboard Models*, in "Symposium on Parallelism in Algorithms and Architectures", United States, 2012, 7, http://hal.inria.fr/hal-00704200.

[26] N. BONICHON, C. GAVOILLE, N. HANUSSE, L. PERKOVIC. *The Stretch Factor of L_1- and L_∞-Delaunay Triangulations*, in "20th Annual European Symposium on Algorithms (ESA)", Ljubljana, Slovenia, September 2012, vol. 7501 of Lecture Notes in Computer Science, p. 205-216, http://hal.inria.fr/hal-00725844.

[27] J. CZYZOWICZ, A. KOSOWSKI, A. PELC. *Time vs. space trade-offs for rendezvous in trees*, in "SPAA - 24th Symposium on Parallelism in Algorithms and Architectures - 2012", Pittsburgh, United States, ACM, 2012, p. 1-10 [*DOI :* 10.1145/2312005.2312007], http://hal.inria.fr/hal-00725924.

[28] G. D'ANGELO, G. DI STEFANO, R. KLASING, A. NAVARRA. *Gathering of Robots on Anonymous Grids without multiplicity detection*, in "19th International Colloquium on Structural Information and Communication Complexity (SIROCCO 2012)", Reykjavík, Iceland, Lecture Notes in Computer Science, Springer, June 2012, vol. 7355, p. 327-338 [*DOI :* 10.1007/978-3-642-31104-8_28], http://hal.inria.fr/hal-00728988.

[29] L. EYRAUD-DUBOIS, H. LARCHEVÊQUE. *Optimizing Resource allocation while handling SLA violations in Cloud Computing platforms*, in "IPDPS - 27th IEEE International Parallel & Distributed Processing Symposium", Boston, États-Unis, 2013, http://hal.inria.fr/hal-00772846.

[30] C. GAVOILLE, P. HALFTERMEYER. *Connexité dans l'urgence*, in "14èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications (AlgoTel)", La Grande Motte, France, F. MATHIEU, N. HANUSSE (editors), April 2012, p. 1-4, http://hal.inria.fr/hal-00690561.

[31] D. ILCINKAS, A. WADE. *Puissance de l'attente aux stations pour l'exploration des réseaux de transport public*, in "14èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications (AlgoTel)", La Grande Motte, France, F. MATHIEU, N. HANUSSE (editors), 2012, p. 107-110, http://hal.inria.fr/hal-00691015.

[32] A. KOSOWSKI, B. LI, N. NISSE, K. SUCHAN. *k-Chordal Graphs: from Cops and Robber to Compact Routing via Treewidth*, in "14èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications (AlgoTel)", La Grande Motte, France, F. MATHIEU, N. HANUSSE (editors), 2012, ..., http://hal.inria.fr/hal-00687120.

[33] A. KOSOWSKI, B. LI, N. NISSE, K. SUCHAN. *k-Chordal Graphs: from Cops and Robber to Compact Routing via Treewidth*, in "39th International Colloquium on Automata, Languages and Programming (ICALP), track C", United Kingdom, 2012, p. 610-622, http://hal.inria.fr/hal-00704202.

[34] A. KOSOWSKI, A. NAVARRA, D. PAJAK, C. PINOTTI. *Maximum Matching in Multi-Interface Networks*, in "COCOA - 6th Conference on Combinatorial Optimization and Applications - 2012", Banff, Canada, Lecture Notes in Computer Science, Springer, 2012, vol. 7402, p. 13-24 [*DOI :* 10.1007/978-3-642-31770-5_2], http://hal.inria.fr/hal-00725925.

National Conferences with Proceeding

[35] L. EYRAUD-DUBOIS, P. UZNANSKI. *Bedibe: Datasets and Software Tools for Distributed Bandwidth Prediction*, in "AlgoTel'2012 - Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications", La Grande Motte, France, N. HANUSSE, F. MATHIEU (editors), May 2012, http://hal.inria.fr/hal-00690861.

Scientific Books (or Scientific Book chapters)

[36] R. KLASING. *Experimental Algorithms - 11th International Symposium, SEA 2012, Bordeaux, France, June 7-9, 2012. Proceedings*, Springer, 2012, 406, Lecture Notes in Computer Science, Vol. 7276 [*DOI :* 10.1007/978-3-642-30850-5], http://hal.inria.fr/hal-00720762.

Research Reports

[37] N. BONICHON, C. GAVOILLE, N. HANUSSE, L. PERKOVIC. *The Stretch Factor of L_1- and L_∞-Delaunay Triangulations*, February 2012, http://hal.inria.fr/hal-00673187.

[38] A. KOSOWSKI, B. LI, N. NISSE, K. SUCHAN. *k-Chordal Graphs: from Cops and Robber to Compact Routing via Treewidth*, Inria, February 2012, n^{o} RR-7888, http://hal.inria.fr/hal-00671861.

Other Publications

[39] N. BONICHON, C. GAVOILLE, N. HANUSSE, L. PERKOVIC. *The Stretch Factor of L_1- and L_∞-Delaunay Triangulations*, May 2012, Poster at STOC, http://hal.inria.fr/hal-00725842.

[40] A. COLLINS, J. CZYZOWICZ, L. GASIENIEC, A. KOSOWSKI, E. KRANAKIS, D. KRIZANC, R. MARTIN, O. MORALES PONCE. *Optimal Patrolling of Fragmented Boundaries*, http://hal.inria.fr/hal-00686270.

[41] J. CZYZOWICZ, D. DERENIOWSKI, L. GASIENIEC, R. KLASING, A. KOSOWSKI, D. PAJAK. *Collision-Free Network Exploration*, http://hal.inria.fr/hal-00736276.

[42] J. CZYZOWICZ, L. GASIENIEC, A. KOSOWSKI, E. KRANAKIS, O. MORALES PONCE, E. PACHECO. *Position Discovery for a System of Bouncing Robots*, http://hal.inria.fr/hal-00735828.

[43] R. KLASING, A. KOSOWSKI, D. PAJAK, T. SAUERWALD. *The Multi-Agent Rotor-Router on the Ring: A Deterministic Alternative to Parallel Random Walks*, http://hal.inria.fr/hal-00735113.

[44] A. KOSOWSKI. *Faster Walks in Graphs: A $\widetilde{O}(n^2)$ Time-Space Trade-off for Undirected s-t Connectivity*, Version 3 makes use of the Metropolis-Hastings walk, http://hal.inria.fr/hal-00685373.

[45] A. KOSOWSKI, P. UZNANSKI. *Splittable Single Source-Sink Routing on CMP Grids: A Sublinear Number of Paths Suffice*, http://hal.inria.fr/hal-00737611.

References in notes

[46] I. ABRAHAM, S. CHECHIK, C. GAVOILLE. *Fully dynamic approximate distance oracles for planar graphs via forbidden-set distance labels*, in "44th Annual ACM Symposium on Theory of Computing (STOC)", New-York, États-Unis, ACM (editor), 2012, p. 1199-1217, http://hal.archives-ouvertes.fr/hal-00725839.

[47] I. ABRAHAM, S. CHECHIK, C. GAVOILLE, D. PELEG. *Forbidden-Set Distance Labels for Graphs of Bounded Doubling Dimension*, in "29th Annual ACM Symposium on Principles of Distributed Computing (PODC)", ACM Press, July 2010, p. 192-200 [*DOI :* 10.1145/1835698.1835743].

[48] I. ATSONIOS, O. BEAUMONT, N. HANUSSE, Y. KIM. *On Power-Law Distributed Balls in Bins and its Applications to View Size Estimation*, in "ISAAC", Yokohama, Japon, 2011, http://hal.inria.fr/inria-00618785.

[49] E. BAMPAS, J. CZYZOWICZ, L. GASIENIEC, D. ILCINKAS, A. LABOUREL. *Almost Optimal Asynchronous Rendezvous in Infinite Multidimensional Grids*, in "DISC 2010", États-Unis, Springer Berlin / Heidelberg, 2010, vol. 6343, p. 297–311, http://hal.archives-ouvertes.fr/hal-00534263.

[50] E. BAMPAS, L. GASIENIEC, N. HANUSSE, D. ILCINKAS, R. KLASING, A. KOSOWSKI. *Euler Tour Lock-in Problem in the Rotor-Router Model*, in "DISC 2009", Espagne, Springer Berlin / Heidelberg, 2009, vol. 5805, p. 421–433, http://hal.archives-ouvertes.fr/hal-00402749.

[51] E. BAMPAS, L. GASIENIEC, R. KLASING, A. KOSOWSKI, T. RADZIK. *Robustness of the rotor-router mechanism*, in "OPODIS 2009", France, Springer Berlin / Heidelberg, 2009, vol. 5923, p. 345–358, http://hal.archives-ouvertes.fr/hal-00417639.

[52] C. BANINO-ROKKONES, O. BEAUMONT, H. REJEB. *Scheduling Techniques for Effective System Reconfiguration in Distributed Storage Systems*, in "IEEE Intl Conference on Parallel and Distributed Systems Conferences", Melbourne, Australie, 2008, http://hal.inria.fr/inria-00336242.

[53] O. BEAUMONT, N. BONICHON, L. COURTÈS, X. HANIN, E. DOLSTRA. *Mixed Data-Parallel Scheduling for Distributed Continuous Integration*, in "Heterogeneity in Computing Workshop, in IPDPS 2012", Shangaï, Chine, IEEE (editor), 2012, http://hal.inria.fr/hal-00684220.

[54] O. BEAUMONT, N. BONICHON, P. DUCHON, H. LARCHEVÊQUE. *Distributed Approximation Algorithm for Resource Clustering*, in "SIROCCO 2008", Suisse, 2008, p. 61–73, http://hal.archives-ouvertes.fr/hal-00323068.

[55] O. BEAUMONT, N. BONICHON, P. DUCHON, H. LARCHEVÊQUE. *Use of Internet Embedding Tools for Heterogeneous Resources Aggregation*, in "Heterogeneity in Computing Workshop", Anchorage, États-Unis, 2011, 10, http://hal.inria.fr/inria-00562655.

[56] O. BEAUMONT, N. BONICHON, P. DUCHON, H. LARCHEVÊQUE. *Utilisation d'outils de plongement d'Internet pour l'agrégation de ressources hétérogènes*, in "Conférence : 13es Rencontres Francophones sur les Aspects Algorithmiques de Télécommunications (AlgoTel)", Cap Estérel, France, P. DUCOURTHIAL (editor), 2011, http://hal.inria.fr/inria-00585254.

[57] O. BEAUMONT, N. BONICHON, L. EYRAUD-DUBOIS. *Scheduling divisibleworkloads on heterogeneous platforms under bounded multi-port model*, in "Heterogeneity in Computing Workshop, in IEEE International Symposium on Parallel and Distributed Processing, 2008. IPDPS 2008.", Miami, FL, États-Unis, 2008, http://hal.inria.fr/inria-00336189.

[58] O. BEAUMONT, N. BONICHON, L. EYRAUD-DUBOIS, L. MARCHAL. *Minimizing Weighted Mean Completion Time for Malleable Tasks Scheduling*, in "IPDPS 2012, 26th IEEE International Parallel & Distributed Processing Symposium", Shangai, Chine, IEEE (editor), 2012, http://hal.inria.fr/inria-00564056.

[59] O. BEAUMONT, N. BONICHON, L. EYRAUD-DUBOIS, P. UZNANSKI. *Broadcasting on Large Scale Heterogeneous Platforms with connectivity artifacts under the Bounded Multi-Port Model*, in "ICPADS 2011 - 17th Internation Conference on Parallel and Distributed Systems", Taïnan, Taïwan, Province De Chine, IEEE (editor), 2011, vol. Proceeding of ICPADS 2011, p. 173-180, http://hal.inria.fr/inria-00618784.

[60] O. BEAUMONT, N. BONICHON, H. LARCHEVÊQUE. *Modeling and Practical Evaluation of a Service Location Problem in Large Scale Networks*, in "Internation Conference on Parallel Computing 2011", Taipei, Taïwan, Province De Chine, 2011, 10, http://hal.inria.fr/inria-00580969.

[61] O. BEAUMONT, P. DUCHON, M. KORZENIOWSKI. *Heterogenous dating service with application to rumor spreading*, in "IEEE International Symposium on Parallel and Distributed Processing, 2008. IPDPS 2008.", Miami, FL, États-Unis, IEEE, 2008, p. 1–10, http://hal.inria.fr/inria-00142778.

[62] O. BEAUMONT, L. EYRAUD-DUBOIS, A. KUMAR. *Broadcasting on Large Scale Heterogeneous Platforms under the Bounded Multi-Port Model*, in "24th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2010)", Atlanta, États-Unis, 2010, http://hal.inria.fr/inria-00444586.

[63] O. BEAUMONT, L. EYRAUD-DUBOIS, H. REJEB, C. THRAVES. *Allocation of Clients to Multiple Servers on Large Scale Heterogeneous Platforms*, in "The Fifteenth International Conference on Parallel and Distributed Systems (ICPADS'09)", Shenzhen, Chine, 2009, http://hal.inria.fr/inria-00413437.

[64] O. BEAUMONT, L. EYRAUD-DUBOIS, H. REJEB, C. THRAVES. *Online Allocation of Splitable Clients to Multiple Servers on Large Scale Heterogeneous Platforms*, in "AlgoTel", Carry-Le-Rouet, France, A. CHAINTREAU, C. MAGNIEN (editors), 2009, http://hal.inria.fr/inria-00384569.

[65] O. BEAUMONT, L. EYRAUD-DUBOIS, H. REJEB, C. THRAVES. *On-line Allocation of Clients to Multiple Servers on Large Scale Heterogeneous Systems*, in "PDP 2010 - The 18th Euromicro International Conference on Parallel, Distributed and Network-Based Computing", Pisa, Italie, 2010, http://hal.inria.fr/inria-00444584.

[66] O. BEAUMONT, L. EYRAUD-DUBOIS, J. WON. *Using the Last-mile Model as a Distributed Scheme for Available Bandwidth Prediction*, in "EuroPar 2011", Bordeaux, France, SPRINGER (editor), 2011, http://hal.inria.fr/inria-00588651.

[67] O. BEAUMONT, H. REJEB. *On the Importance of Bandwidth Control Mechanisms for Scheduling on Large Scale Heterogeneous Platforms*, in "24th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2010)", Atlanta, États-Unis, 2010, http://hal.inria.fr/inria-00444585.

[68] N. BONICHON, C. GAVOILLE, N. HANUSSE, D. ILCINKAS. *Connections between Theta-Graphs, Delaunay Triangulations, and Orthogonal Surfaces*, in "WG 2010", Grèce, Springer Berlin / Heidelberg, 2010, vol. 6410, p. 266–278, http://hal.archives-ouvertes.fr/hal-00536710.

[69] N. BONICHON, C. GAVOILLE, N. HANUSSE, L. PERKOVIC. *Plane Spanners of Maximum Degree Six*, in "ICALP", France, 2010, vol. 6198, p. 19-30, http://hal.archives-ouvertes.fr/hal-00534212.

[70] N. BONICHON, C. GAVOILLE, N. HANUSSE, L. PERKOVIĆ. *The Stretch Factor of L_1- and L_∞-Delaunay Triangulations*, in "20th Annual European Symposium on Algorithms (ESA)", L. EPSTEIN, P. FERRAGINA (editors), Springer-Verlag, September 2012, vol. 7501 of Lecture Notes in Computer Science, p. 205-216.

[71] J. CHALOPIN, S. DAS, A. KOSOWSKI. *Constructing a Map of an Anonymous Graph: Applications of Universal Sequences*, in "OPODIS - 14th International Conference On Principles Of DIstributed Systems - 2010", Tozeur, Tunisie, Springer, 2010, vol. 6490, p. 119-134, http://hal.inria.fr/hal-00726071.

[72] B. CHUN, D. CULLER, T. ROSCOE, A. BAVIER, L. PETERSON, M. WAWRZONIAK, M. BOWMAN. *PlanetLab: an overlay testbed for broad-coverage services*, in "ACM SIGCOMM Comput. Commun. Rev.", July 2003, vol. 33, p. 3–12, http://doi.acm.org/10.1145/956993.956995.

[73] R. COHEN, P. FRAIGNIAUD, D. ILCINKAS, A. KORMAN, D. PELEG. *Label-Guided Graph Exploration by a Finite Automaton*, in "ACM Transactions on Algorithms", 2008, vol. 4, n^o 4, Article 42, http://hal.archives-ouvertes.fr/hal-00341609.

[74] A. COLLINS, J. CZYZOWICZ, L. GASIENIEC, A. KOSOWSKI, A. RUSSELL. *Synchronous Rendezvous for Location-Aware Agents*, in "DISC 2011 - 25th International Symposium on Distributed Computing", Rome, Italie, D. PELEG (editor), Springer, 2011, vol. 6950, p. 447-459, http://hal.inria.fr/hal-00646907.

[75] C. COOPER, D. ILCINKAS, R. KLASING, A. KOSOWSKI. *Derandomizing Random Walks in Undirected Graphs Using Locally Fair Exploration Strategies*, in "ICALP 2009", Grèce, Springer Berlin / Heidelberg, 2009, vol. 5556, p. 411-422, http://hal.archives-ouvertes.fr/hal-00374071.

[76] C. COOPER, R. KLASING, T. RADZIK. *Locating and repairing faults in a network with mobile agents*, in "Proceedings of the 15th Colloquium on Structural Information and Communication Complexity (SIROCCO 2008)", Suisse, 2008, vol. 5058, p. 20–32, http://hal.archives-ouvertes.fr/hal-00307025.

[77] J. CZYZOWICZ, S. DOBREV, L. GASIENIEC, D. ILCINKAS, J. JANSSON, R. KLASING, I. LIGNOS, R. MARTIN, K. SADAKANE, W.-K. SUNG. *More Efficient Periodic Traversal in Anonymous Undirected Graphs*, in "SIROCCO 2009", Slovénie, Springer Berlin / Heidelberg, 2009, vol. 5869, p. 174–188, http://hal.archives-ouvertes.fr/hal-00371489.

[78] J. CZYZOWICZ, L. GASIENIEC, A. KOSOWSKI, E. KRANAKIS. *Boundary Patrolling by Mobile Agents with Distinct Maximal Speeds*, in "ESA 2011 - 19th Annual European Symposium on Algorithms", Saarbruecken, Allemagne, C. DEMETRESCU, M. M. HALLDÓRSSON (editors), Springer, 2011, vol. 6942, p. 701-712, http://hal.inria.fr/hal-00646910.

[79] J. CZYZOWICZ, D. ILCINKAS, A. LABOUREL, A. PELC. *Asynchronous Deterministic Rendezvous in Bounded Terrains*, in "SIROCCO 2010", Şirince, Turquie, Springer Berlin / Heidelberg, 2010, vol. 6058, p. 72-85, http://hal.archives-ouvertes.fr/hal-00516048.

[80] J. CZYZOWICZ, D. ILCINKAS, A. LABOUREL, A. PELC. *Optimal Exploration of Terrains with Obstacles*, in "SWAT 2010", Bergen, Norvège, Springer Berlin / Heidelberg, 2010, vol. 6139, p. 1-12, http://hal.archives-ouvertes.fr/hal-00516061.

[81] J. CZYZOWICZ, A. KOSOWSKI, A. PELC. *How to meet when you forget: log-space rendezvous in arbitrary graphs*, in "Distributed Computing", 2012, vol. 25, n^o 2, p. 165-178, http://hal.inria.fr/hal-00725875.

[82] J. CZYZOWICZ, A. KOSOWSKI, A. PELC. *Time vs. space trade-offs for rendezvous in trees*, in "SPAA - 24th Symposium on Parallelism in Algorithms and Architectures - 2012", Pittsburgh, États-Unis, 2012, p. 1-10, http://hal.inria.fr/hal-00725924.

[83] F. DABEK, R. COX, F. KAASHOEK, R. MORRIS. *Vivaldi: a decentralized network coordinate system*, in "SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications", New York, NY, USA, ACM, 2004, p. 15–26, http://doi.acm.org/10.1145/1015467.1015471.

[84] E. DIOT, C. GAVOILLE. *On the Path Separability of Planar Graphs*, in "EuroComb 2009, Bordeaux, France", Bordeaux, France, 2009, vol. 34C, p. 549-552, http://hal.archives-ouvertes.fr/hal-00408228.

[85] P. DUCHON, H. LARCHEVÊQUE. *On the search path length of random binary skip graphs*, in "ANALCO 2010", Austin, États-Unis, SIAM (editor), SIAM, 2010, p. 1-8, http://hal.inria.fr/inria-00547935.

[86] L. EYRAUD-DUBOIS, P. UZNANSKI. *Bedibe: Datasets and Software Tools for Distributed Bandwidth Prediction*, in "AlgoTel'2012 - Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications", La Grande Motte, France, N. HANUSSE, F. MATHIEU (editors), 2012, http://hal.inria.fr/hal-00690861.

[87] P. FLOCCHINI, D. ILCINKAS, A. PELC, N. SANTORO. *Remembering without Memory: Tree Exploration by Asynchronous Oblivious Robots*, in "SIROCCO 2008", Villars-sur-Ollon, Suisse, A. A. SHVARTSMAN, P. FELBER (editors), Springer Berlin / Heidelberg, 2008, vol. 5058, p. 33-47, http://hal.archives-ouvertes.fr/hal-00341465.

[88] P. FLOCCHINI, D. ILCINKAS, A. PELC, N. SANTORO. *Remembering Without Memory: Tree Exploration by Asynchronous Oblivious Robots*, in "Theoretical Computer Science", 2010, vol. 411, n^{o} 14-15, p. 1583-1598, http://hal.archives-ouvertes.fr/hal-00464588.

[89] P. FLOCCHINI, D. ILCINKAS, N. SANTORO. *Ping Pong in Dangerous Graphs: Optimal Black Hole Search with Pure Tokens*, in "DISC 2008", France, G. TAUBENFELD (editor), Springer Berlin / Heidelberg, 2008, vol. 5218, p. 227-241, http://hal.archives-ouvertes.fr/hal-00341523.

[90] C. GAVOILLE, S. CHRISITAN. *Sparse Spanners vs. Compact Routing*, in "23rd Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)", San Jose, États-Unis, ACM, 2011, p. 225-234, http://hal.archives-ouvertes.fr/hal-00651836.

[91] C. GAVOILLE, Q. GODFROY, L. VIENNOT. *Node-Disjoint Multipath Spanners and their Relationship with Fault-Tolerant Spanners*, in "OPODIS'11 - 15th International Conference on Principles of Distributed Systems", Toulouse, France, A. F. ANTA, G. LIPARI, M. ROY (editors), Springer, 2011, vol. 7109, p. 143-158, http://hal.archives-ouvertes.fr/hal-00651825.

[92] C. GAVOILLE, D. PELEG, S. PÉRENNES, R. RAZ. *Distance labeling in graphs*, in "Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms", Society for Industrial and Applied Mathematics, 2001, 219.

[93] N. HANUSSE, D. ILCINKAS, A. KOSOWSKI, N. NISSE. *Locating a Target with an Agent Guided by Unreliable Local Advice*, in "PODC 2010", Zurich, Suisse, ACM New York, NY, USA, 2010, p. 355-364, http://hal.archives-ouvertes.fr/hal-00516695.

[94] N. HANUSSE, D. KAVVADIAS, E. KRANAKIS, D. KRIZANC. *Memoryless search algorithms in a network with faulty advice*, in "Theoretical Computer Science", 2008, vol. 402, n^{o} 2-3, p. 190-198, http://hal.archives-ouvertes.fr/hal-00341468.

[95] N. HANUSSE, S. MAABOUT. *A parallel Algorithm for Computing Borders*, in "International Conference on Information and Knowledge Managment (CIKM)", Glasgow, Royaume-Uni, ACM, 2011, p. 1639-1648, http://hal.archives-ouvertes.fr/hal-00644517.

[96] N. HANUSSE, S. MAABOUT, R. TOFAN. *A view selection algorithm with performance guarantee*, in "International Conference on Extending Database Technology", St Petersburg, Russie, Fédération De, 2009, électronique, http://hal.archives-ouvertes.fr/hal-00341475.

[97] N. HANUSSE, S. MAABOUT, R. TOFAN. *Revisiting the Partial Data Cube Materialization*, in "Conference on Advances in Data Bases and Information Systems (ADBIS)", Vienne, Autriche, Springer, 2011, n^o 6909, p. 70-83, http://hal.archives-ouvertes.fr/hal-00644526.

[98] R. KLASING, A. KOSOWSKI, A. NAVARRA. *Taking Advantage of Symmetries: Gathering of many Asynchronous Oblivious Robots on a Ring*, in "Theoretical Computer Science", 2010, vol. 411, n^o 34–36, p. 3235–3246, http://hal.archives-ouvertes.fr/hal-00519069.

[99] R. KLASING, E. MARKOU, A. PELC. *Gathering asynchronous oblivious mobile robots in a ring*, in "Theoretical Computer Science", 2008, vol. 390, n^o 1, p. 27–39, http://hal.archives-ouvertes.fr/hal-00307237.

[100] A. KOSOWSKI, B. LI, N. NISSE, K. SUCHAN. *k-Chordal Graphs: from Cops and Robber to Compact Routing via Treewidth*, in "14èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications (AlgoTel)", La Grande Motte, France, N. MATHIEU (editor), 2012, ..., http://hal.archives-ouvertes.fr/hal-00687120.

[101] A. KOSOWSKI, A. NAVARRA. *Graph Decomposition for Memoryless Periodic Exploration*, in "Algorithmica", 2012, vol. 63, n^o 1-2, p. 26-38, http://hal.inria.fr/hal-00725922.

[102] A. KOSOWSKI, A. NAVARRA, C. PINOTTI. *Synchronous Black Hole Search in Directed Graphs*, in "Theoretical Computer Science", 2011, vol. 412, n^o 41, p. 5752-5759, http://hal.inria.fr/inria-00614476.

[103] A. LEGRAND, F. MAZOIT, M. QUINSON. *An Application-level network mapper*, LIP, ENS Lyon, France, feb 2003, n^o RR-2003-09.

[104] Y. MAO, L. SAUL, J. SMITH. *IDES: An Internet Distance Estimation Service for Large Networks*, in "IEEE JSAC", Dec. 2006, vol. 24, n^o 12, p. 2273 -2284.

[105] J.-F. MARCKERT, N. BONICHON. *Asymptotics of geometrical navigation on a random set of points of the plane*, in "Advances in Applied Probability", 2011, vol. 43, n^o 4, p. 899-942, http://hal.inria.fr/hal-00649901.

[106] V. RAMASUBRAMANIAN, D. MALKHI, F. KUHN, M. BALAKRISHNAN, A. GUPTA, A. AKELLA. *On the treeness of internet latency and bandwidth*, in "SIGMETRICS '09: Proceedings of the eleventh international joint conference on Measurement and modeling of computer systems", New York, NY, USA, ACM, 2009, p. 61–72.

[107] H. WATANABE, H. NAKAZATO, A. GHATPANDE, O. BEAUMONT. *Divisible Load Scheduling with Result Collection on Heterogeneous Systems*, in "IEEE International Symposium on Parallel and Distributed Processing, 2008. IPDPS 2008.", Miami, FL, États-Unis, 2008, http://hal.inria.fr/inria-00336186.