



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team gang

Networks, Graphs and Algorithms

Paris - Rocquencourt

Theme : Networks and Telecommunications

Activity
R *eport*

2010

Table of contents

1. Team	1
2. Overall Objectives	1
3. Application Domains	2
4. New Results	2
4.1. Understanding graph representations	2
4.1.1. Efficient graph spanners	2
4.1.2. Web graph algorithms	2
4.1.3. Modular decomposition	2
4.1.4. Unique Perfect phylogeny tree is NP-hard	2
4.2. Small world networks structure	2
4.2.1. Small world searchability	2
4.2.2. Rumor spreading complexity	3
4.2.3. Structured/Unstructured Overlays	3
4.3. Distributed Online Algorithms	3
4.3.1. Efficient Threshold Detection in a Distributed Environment	3
4.3.2. On the Additive Constant of the k -Server Work Function Algorithm	3
4.4. Informative Labeling Schemes	3
4.4.1. Optimal Ancestry Scheme and small Universal Posets	3
4.4.2. Compact Ancestry Labeling Schemes for XML Trees	4
4.4.3. Labeling Schemes for Vertex Connectivity	4
4.4.4. Constructing Labeling Schemes through Universal Matrices	4
4.5. Proof Labels	5
4.6. Fault Tolerance in Distributed Networks	5
4.6.1. Weakest failure Detection	5
4.6.2. Fault Tolerance and Stabilization	5
4.6.3. Byzantine agreement	6
4.7. Discrete Optimization Algorithms	6
4.7.1. Non Uniform Selection of Solutions for Upper Bounding the 3-SAT Threshold	6
4.7.2. Estimating Satisfiability	6
4.7.3. Eigenvectors of Toeplitz matrices under higher order three term recurrence and circulant perturbation	7
4.7.4. Convex Maximization problems	7
5. Other Grants and Activities	7
5.1. Regional Initiatives	7
5.2. National initiatives	7
5.2.1. ANR Algorithm Design and Analysis for Implicitly and Incompletely Defined Interaction Networks (ALADIN)	7
5.2.2. ANR PROSE	7
5.2.3. ANR Shaman	7
5.3. European Initiatives	8
6. Dissemination	8
6.1. Animation of the scientific community	8
6.2. Teaching	8
6.3. Theses	9
6.3.1. Defended theses	9
6.3.2. Ongoing theses	9
6.3.3. PostDoc	9
7. Bibliography	9

1. Team

Research Scientists

Laurent Viennot [Team leader, Research Associate (CR Inria), HdR]
Dominique Fortin [Research Associate (CR Inria)]
Pierre Fraigniaud [Research Director (DR Cnrs), HdR]
Amos Korman [Research Associate (CR Cnrs)]
Fabien Mathieu [Research Associate]

Faculty Members

Michel Habib [Professor (Paris Diderot Univ.), HdR]
Yacine Boufkhad [Assistant professor (Paris Diderot Univ.)]
Pierre Charbit [Assistant professor (Paris Diderot Univ.)]
Hugues Fauconnier [Professor (Paris Diderot Univ.)]
Carole Gallet-Delporte [Professor (Paris Diderot Univ.)]
Fabien de Montgolfier [Assistant professor (Paris Diderot Univ.)]

PhD Students

Thomas Hugel [PhD]
Hoang Anh Phan [PhD]
Heger Arfaoui [PhD Student]
Hervé Baumann [PhD Student]
Xavier Koegler [PhD Student]
Antoine Mamcarz [PhD Student]
Mauricio Soto [PhD Student]

Post-Doctoral Fellows

George Giakkoupis [PostDoc]
Julien Clément [PostDoc]

Administrative Assistant

Christine Anocq [shared time (with Hipercom)]

2. Overall Objectives

2.1. Overall Objectives

Our goal is to develop the field of graph algorithms for networks. Based on algorithmic graph theory and graph modeling we want to understand what can be done in these large networks and what cannot. Furthermore, we want to derive practical distributed algorithms from known strong theoretical results. Finally, we want to extract possibly new graph problems by focusing on particular applications.

The main goal to achieve in networks are efficient searching of nodes or data, and efficient content transfers. We propose to implement strong theoretical results in that domain to make significant breakthrough in large network algorithms. These results concern small world routing, low stretch routing in doubling metrics and bounded width classes of graphs. They are detailed in the next section. This implies several challenges:

- testing our target networks against general graph parameters known to bring theoretical tractability,
- implementing strong theoretical results in the dynamic and distributed context of large networks.

A complementary approach consists in studying the combinatorial and graph structures that appear in our target networks. These structures may have inherent characteristics coming from the way the network is formed, or from the design goals of the target application.

3. Application Domains

3.1. Application Domains

Application domains include evaluating Internet performances, the design of new peer-to-peer applications, enabling large scale ad hoc networks and mapping the web.

- The application of measuring and modeling Internet metrics such as latencies and bandwidth is to provide tools for optimizing Internet applications. This concerns especially large scale applications such as web site mirroring and peer-to-peer applications.
- Peer-to-peer protocols are based on a all equal paradigm that allows to design highly reliable and scalable applications. Besides the file sharing application, peer-to-peer solutions could take over in web content dissemination resistant to high demand bursts or in mobility management. Envisioned peer-to-peer applications include video on demand, streaming, exchange of classified ads,...
- Wifi networks have entered our every day life. However, enabling them at large scale is still a challenge. Algorithmic breakthrough in large ad hoc networks would allow to use them in fast and economic deployment of new radio communication systems.
- The main application of the web graph structure consists in ranking pages. Enabling site level indexing and ranking is a possible application o f such studies.

4. New Results

4.1. Understanding graph representations

4.1.1. *Efficient graph spanners*

Participants: Cyril Gavoille [CNRS LABRI, University of Bordeaux, France], Quentin Godfroy [ENS, Paris, France], Laurent Viennot.

In [26], we investigate the construction of sparse spanners preserving multiple paths.

4.1.2. *Web graph algorithms*

Participants: Fabien Mathieu, Laurent Viennot.

A result from 2006 about page ranking versus site partition of the web is published this year [28].

4.1.3. *Modular decomposition*

Participants: Michel Habib, Christophe Paul.

In [10], we review algorithmic aspects of modular decomposition.

4.1.4. *Unique Perfect phylogeny tree is NP-hard*

Participant: Michel Habib.

In [34], we prove an old conjecture about uniqueness of perfect phylogeny tree.

4.2. Small world networks structure

4.2.1. *Small world searchability*

Participants: Pierre Fraigniaud, George Giakkoupis.

In [18], we analyze decentralized routing in small-world networks with an arbitrary underlying structure.

4.2.2. Rumor spreading complexity

Participants: Pierre Fraigniaud, George Giakkoupis.

In [17] we analyze the bit communication complexity of randomized rumor spreading.

4.2.3. Structured/Unstructured Overlays

Participants: Pierre Fraigniaud, Hoang Anh Phan.

On the one hand, in [24], we study the impact of degree balancing in de Bruijn-Based Overlay Networks; in the other hand, in [23] we design Tree-Based Multicast Schemes in Peer-to-Peer Overlay Networks.

4.3. Distributed Online Algorithms

4.3.1. Efficient Threshold Detection in a Distributed Environment

Participants: Yuval Emek [University of Tel Aviv, Israel], Amos Korman [CNRS LIAFA, University of Paris Diderot, France].

Consider a distributed network in which *events* occur at arbitrary nodes and at unpredictable times. An event occurring at node u is sensed only by u which in turn may invoke a communication protocol that allows nodes to exchange messages with their neighbors. We are interested in the following *threshold detection (TD)* problem inherent to distributed computing: Given some threshold k , the goal of a TD protocol is to broadcast a *termination* signal when at least k events have occurred (throughout the network).

In [16], we develop a randomized TD protocol that may fail with negligible probability but which significantly improves previous results in terms of the message complexity, namely, the total number of messages sent by all participating nodes. With the right choice of parameters our randomized protocol turns into a deterministic one that guarantees low communication burden for any node. This is a principal complexity measure in many applications of wireless networks and which, to the best of our knowledge, has not been bounded before in the context of such problems.

4.3.2. On the Additive Constant of the k -Server Work Function Algorithm

Participants: Yuval Emek [University of Tel Aviv, Israel], Pierre Fraigniaud [CNRS LIAFA, University of Paris Diderot, France], Amos Korman [CNRS LIAFA, University of Paris Diderot, France], Adi Rosen [CNRS LRI, University of Paris Sud, France].

In [3], we consider the Work Function Algorithm for the k -server problem. We show that if the Work Function algorithm is c -competitive, then it is also strictly $2c$ -competitive. Thus, as a consequence of a famous result by Koutsoupias and Papadimitriou (J. ACM 95) we get that the Work Function algorithm is strictly $(4k - 2)$ -competitive.

4.4. Informative Labeling Schemes

4.4.1. Optimal Ancestry Scheme and small Universal Posets

Participants: Pierre Fraigniaud [CNRS LIAFA, University of Paris Diderot, France], Amos Korman [CNRS LIAFA, University of Paris Diderot, France].

In [19], we solve the *ancestry problem*, which was introduced more than twenty years ago by Kannan et al. [STOC '88], and is among the most well-studied problems in the field of informative labeling schemes. Specifically, we construct an *ancestry labeling scheme* for n -node trees with label size $\log_2 n + O(\log \log n)$ bits, thus matching the $\log_2 n + \Omega(\log \log n)$ bits lower bound given by Alstrup et al. [SODA '03]. Besides its optimal label size, our scheme assigns the labels in linear time, and guarantees that any ancestry query can be answered in constant time.

In addition to its potential impact in terms of improving the performances of XML search engines, our ancestry scheme is also useful in the context of partially ordered sets. Specifically, for any fixed integer k , our scheme enables the construction of a *universal* poset of size $O(n^k \log^{4k} n)$ for the family of n -element posets with tree-dimension at most k . This bound is almost tight thanks to a lower bound of $n^{k-o(1)}$ due to Alon and Scheinerman [Order '88].

4.4.2. Compact Ancestry Labeling Schemes for XML Trees

Participants: Pierre Fraigniaud [CNRS LIAFA, University of Paris Diderot, France], Amos Korman [CNRS LIAFA, University of Paris Diderot, France].

An *ancestry labeling scheme* labels the nodes of any tree in such a way that ancestry queries between any two nodes can be answered just by looking at their corresponding labels. The common measure to evaluate the quality of an ancestry scheme is by its *label size*, that is the maximum number of bits stored in a label, taken over all n -node trees. The design of ancestry labeling schemes finds applications in XML search engines. In these contexts, even small improvements in the label size are important. As a result, following the proposal of a simple interval based ancestry scheme with label size $2 \log n$ bits (Kannan et al., STOC 88), a considerable amount of work was devoted to improve the bound on the label size. The current state of the art upper bound is $\log n + O(\sqrt{\log n})$ bits (Abiteboul et al., SICOMP 06) which is still far from the known $\log n + \Omega(\log \log n)$ lower bound (Alstrup et al., SODA 03). Motivated by the fact that typical XML trees have extremely small depth, [21] parameterizes the quality measure of an ancestry scheme not only by the number of nodes in the given tree but also by its depth. Our main result is the construction of an ancestry scheme that labels n -node trees of depth d with labels of size $\log n + 2 \log d + O(1)$. In addition to our main result, we prove a result that may be of independent interest concerning the existence of a small *universal graph* for the family of trees with bounded depth.

4.4.3. Labeling Schemes for Vertex Connectivity

Participant: Amos Korman [CNRS LIAFA, University of Paris Diderot, France].

In [11] we study labeling schemes for the vertex connectivity function on general graphs. We consider the problem of assigning short labels to the nodes of any n -node graph in such a way that given the labels of any two nodes u and v , one can decide whether u and v are k -vertex connected in G , i.e., whether there exist k vertex disjoint paths connecting u and v . The paper establishes an upper bound of $k^2 \log n$ on the number of bits used in a label. The best previous upper bound for the label size of such a labeling scheme is $2^k \log n$.

4.4.4. Constructing Labeling Schemes through Universal Matrices

Participants: Amos Korman [CNRS LIAFA, University of Paris Diderot, France], David Peleg [Weizmann Institute of Science, Israel], Yoav Rodeh [University of Tel Aviv, Israel].

Let f be a function on pairs of vertices. An f -*labeling scheme* for a family of graphs \mathcal{F} labels the vertices of all graphs in \mathcal{F} such that for every graph $G \in \mathcal{F}$ and every two vertices $u, v \in G$, $f(u, v)$ can be inferred by merely inspecting the labels of u and v . The *size* of a labeling scheme is the maximum number of bits used in a label of any vertex in any graph in \mathcal{F} . In [12] we illustrate that the notion of universal matrices can be used to efficiently construct f -labeling schemes.

Let $\mathcal{F}(n)$ be a family of connected graphs of size at most n and let $\mathcal{C}(\mathcal{F}, n)$ denote the collection of graphs of size at most n , such that each graph in $\mathcal{C}(\mathcal{F}, n)$ is composed of a disjoint union of some graphs in $\mathcal{F}(n)$. We first investigate methods for translating f -labeling schemes for $\mathcal{F}(n)$ to f -labeling schemes for $\mathcal{C}(\mathcal{F}, n)$. In particular, we show that in many cases, given an f -labeling scheme of size $g(n)$ for a graph family $\mathcal{F}(n)$, one can construct an f -labeling scheme of size $g(n) + \log \log n + O(1)$ for $\mathcal{C}(\mathcal{F}, n)$. We also show that in several cases, the above mentioned extra additive term of $\log \log n + O(1)$ is necessary. In addition, we show that the family of n -node graphs which are unions of disjoint circles enjoys an adjacency labeling scheme of size $\log n + O(1)$. This illustrates a non-trivial example showing that the above mentioned extra additive term is sometimes not necessary.

We then turn to investigate distance labeling schemes on the class of circles of at most n vertices and show an upper bound of $1.5 \log n + O(1)$ and a lower bound of $4/3 \log n - O(1)$ for the size of any such labeling scheme.

4.5. Proof Labels

4.5.1. Proof labeling schemes

Participants: Amos Korman [CNRS LIAFA, University of Paris Diderot, France], Shay Kutten [Technion, Israel], David Peleg [Weizmann Institute of Science, Israel].

In [27] we address the problem of locally verifying global properties. Several natural questions are studied, such as “how expensive is local verification?” and more specifically, “how expensive is local verification compared to computation?” A suitable model is introduced in which these questions are studied in terms of the number of bits a vertex needs to communicate. The model includes the definition of a proof labeling scheme (a pair of algorithms- one to assign the labels, and one to use them to verify that the global property holds). In addition, approaches are presented for the efficient construction of schemes, and upper and lower bounds are established on the bit complexity of schemes for multiple basic problems. The paper also studies the role and cost of unique identities in terms of impossibility and complexity, in the context of proof labeling schemes.

Previous studies on related questions deal with distributed algorithms that simultaneously compute a configuration and verify that this configuration has a certain desired property. It turns out that this combined approach enables the verification to be less costly sometimes, since the configuration is typically generated so as to be easily verifiable. In contrast, our approach separates the configuration design from the verification. That is, it first generates the desired configuration without bothering with the need to verify it, and then handles the task of constructing a suitable verification scheme. Our approach thus allows for a more modular design of algorithms, and has the potential to aid in verifying properties even when the original design of the structures for maintaining them was done without verification in mind.

4.6. Fault Tolerance in Distributed Networks

Participants: Hugues Fauconnier, Carole Gallet-Delporte.

4.6.1. Weakest failure Detection

In [2], we determine the weakest failure detectors to implement shared atomic objects in a distributed system with crash-prone processes. We first determine the weakest failure detector for the basic register object. We then use that to determine the weakest failure detector for all popular atomic objects including test-and-set, fetch-and-add, queue, consensus and compare-and-swap, which we show is the same.

4.6.2. Fault Tolerance and Stabilization

In [1] This article deals with stabilization and fault-tolerance. We consider two types of stabilization: the self- and the pseudo- stabilization. Our goal is to implement the self- and/or pseudo- stabilizing leader election in systems with process crashes, weak reliability, and synchrony assumptions. We try to propose, when it is possible, communication efficient implementations. Our approach allows to obtain algorithms that tolerate both transient and crash failures. Note that some of our solutions are adapted from existing fault-tolerant algorithms. The motivation here is not to propose new algorithms but merely to show some assumptions required to obtain stabilizing leader elections in systems with crash failures. In particular, we focus on the borderline assumptions where we go from the possibility to have self-stabilizing solutions to the possibility to only have pseudo-stabilizing ones.

We consider asynchronous message-passing systems in which some links are timely and processes may crash. Each run defines a timeliness graph among correct processes: (p, q) is an edge of the timeliness graph if the link from p to q is timely (that is, there is a bound on communication delays from p to q). The main goal in [14] is to approximate this timeliness graph by graphs having some properties (such as being trees, rings, ...). Given a family S of graphs, for runs such that the timeliness graph contains at least one graph in S then using an extraction algorithm, each correct process has to converge to the same graph in S that is, in a precise sense, an approximation of the timeliness graph of the run. For example, if the timeliness graph contains a ring, then using an extraction algorithm, all correct processes eventually converge to the same ring and in this ring all nodes will be correct processes and all links will be timely. We first present a general extraction algorithm and then a more specific extraction algorithm that is communication efficient (i.e., eventually all the messages of the extraction algorithm use only links of the extracted graph).

In [13], we study here in which way processes may approximate and find structural properties based on -timeliness (e.g., find timely paths between processes or build a ring between correct processes using only timely links). To that end, we define a notion of approximation of predicates. Then, with help of such approximations, we give a general algorithm that enables to choose and eventually agree on one of these predicates. Finally, applying this approach to timeliness, we give conditions and algorithms to approximate timeliness and dynamically find structural properties using timeliness.

4.6.3. Byzantine agreement

In [15], we address Byzantine agreement in a message passing system with homonyms, i.e. a system with a number l of authenticated identities that is independent of the total number of processes n , in the presence of $t < n$ Byzantine processes. We prove the following results:

- (i) agreement is possible if (and only if) $l > 3t$ in a synchronous model;
- (ii) agreement is impossible, independently of the number of failures, in an eventually synchronous model;
- (iii) eventual agreement is possible, if (and only if) $l > 3t$, in an asynchronous model.

4.7. Discrete Optimization Algorithms

4.7.1. Non Uniform Selection of Solutions for Upper Bounding the 3-SAT Threshold

Participants: Yacine Boufkhad, Thomas Hugel.

In [29], we give a new insight into the upper bounding of the 3-SAT threshold by the first moment method. The best criteria developed so far to select the solutions to be counted discriminate among neighboring solutions on the basis of uniform information about each individual free variable. What we mean by uniform information, is information which does not depend on the solution: e.g. the number of positive/negative occurrences of the considered variable. What is new in our approach is that we use non uniform information about variables. Thus we are able to make a more precise tuning, resulting in a slight improvement on upper bounding the 3-SAT threshold for various models of formulas defined by their distributions.

4.7.2. Estimating Satisfiability

Participants: Yacine Boufkhad, Thomas Hugel.

The problem of estimating the proportion of satisfiable instances of a given CSP (constraint satisfaction problem) can be tackled through two different ways: ordering and weighting. Ordering consists in putting a total order on the domain, which induces an orientation between neighboring solutions in a way that prevents circuits from appearing, and then counting only minimal elements. Weighting consists in putting onto each solution a non-negative real value based on its neighborhood in a way that the total weight is at least 1 for each satisfiable instance. In [33], we investigate the combinatorial properties of these two systems of estimation. First we give some sufficient conditions for a weighting system to be correct, and then we compare it to ordering under different conditions.

4.7.3. Eigenvectors of Toeplitz matrices under higher order three term recurrence and circulant perturbation

Participant: Dominique Fortin.

For Toeplitz matrices following a higher order three term recurrence between the main diagonal and 2 other diagonals l and r apart, and possibly circulant perturbations, we extend the kernel method for computing analytically eigenvectors along their eigenvalues. Unlike, the tridiagonal case with corner perturbations which involves sine functions, it is related to the $(p + 1, p)$ Pascal triangle for the special cases where the diagonals are respectively at distance l and $r = pl$ from the main diagonal [4]

4.7.4. Convex Maximization problems

Participants: Dominique Fortin, Ider Tseveendorj.

In [5] we provide an algorithm, where to escape from a local maximum y of convex function f over D , we (locally) solve piecewise convex maximization $\max \{ \min \{ f(x) \hat{=} f(y), py(x) \} \mid x \hat{=} D \}$ with an additional convex function $py(\hat{\cdot})$. The last problem can be seen as a strictly convex improvement of the standard cutting plane technique for convex maximization. We report some computational results, that show the algorithm efficiency.

5. Other Grants and Activities

5.1. Regional Initiatives

5.1.1. PEFICAMO

Participants: Hugues Fauconnier, Carole Gallet-Delporte, Julien Clément.

Managed by University Paris Diderot, H. Fauconnier is leading this project granting J. Clément from Région Ile de France.

5.2. National initiatives

5.2.1. ANR Algorithm Design and Analysis for Implicitly and Incompletely Defined Interaction Networks (ALADIN)

Participants: Cyril Gavaille [CNRS LABRI, University of Bordeaux, France], Dominique Fortin, Laurent Viennot, Michel Habib, Pierre Charbit, Pierre Fraigniaud.

Pierre Fraigniaud is leading an 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).

5.2.2. ANR PROSE

Participant: Pierre Fraigniaud.

Managed by University Paris Diderot, P. Fraigniaud leads this project.

5.2.3. ANR Shaman

Participants: Hugues Fauconnier, Carole Gallet-Delporte, Hung Tran.

Managed by University Paris Diderot, H. Fauconnier leads this project that grants H. Tran.

5.3. European Initiatives

5.3.1. EULER (*Experimental UpdateLess Evolutive Routing*)

EULER is a 3-year STREP Project targeting Challenge 1 "Technologies and systems architectures for the Future Internet" of the European Commission (EC) Seventh Framework Programme (FP7). The project scope and methodology position within the FIRE (Future Internet Research and Experimentation) Objective ICT-2009.1.6 Part b: Future Internet experimentally-driven research .

The main objective of the EULER exploratory research project is to investigate new routing paradigms so as to design, develop, and validate experimentally a distributed and dynamic routing scheme suitable for the future Internet and its evolution. The resulting routing scheme(s) is/are intended to address the fundamental limits of current stretch-1 shortest-path routing in terms of routing table scalability but also topology and policy dynamics (perform efficiently under dynamic network conditions). Therefore, this project will investigate trade-offs between routing table size (to enhance scalability), routing scheme stretch (to ensure routing quality) and communication cost (to efficiently and timely react to various failures). The driving idea of this research project is to make use of the structural and statistical properties of the Internet topology (some of which are hidden) as well as the stability and convergence properties of the Internet policy in order to specialize the design of a distributed routing scheme known to perform efficiently under dynamic network and policy conditions when these properties are met. The project will develop new models and tools to exhaustively analyse the Internet topology, to accurately and reliably measure its properties, and to precisely characterize its evolution. These models, that will better reflect the network and its policy dynamics, will be used to derive useful properties and metrics for the routing schemes and provide relevant experimental scenarios. The project will develop appropriate tools to evaluate the performance of the proposed routing schemes on large-scale topologies (order of 10k nodes). Prototype of the routing protocols as well as their functional validation and performance benchmarking on the iLAB experimental facility and/or virtual experimental facilities such as PlanetLab/OneLab will allow validating under realistic conditions the overall behaviour of the proposed routing schemes.

More information at <http://www.euler-fire-project.eu/>

6. Dissemination

6.1. Animation of the scientific community

- Laurent Viennot is a scientific editor of the *Interstices* (<http://interstices.info/>) vulgarization site initiated by Inria in collaboration with french universities and Cnrs. He has written an article on the differences between the web and internet.
- Michel Habib is member of the steering committee of STACS (Symposium on Theoretical Aspects of Computer Science) and also of WG (International Workshop on Graph-Theoretic Concepts in Computer Science).

6.2. Teaching

Master MPRI Michel Habib is in charge of a course entitled "graph algorithms". Pierre Fraigniaud is in charge of the course "Algorithmique distribuée pour les réseaux". Laurent Viennot is teaching "Networks and geometry".

D.U.T., University of Paris Diderot Yacine Boufkhad is teaching scientific computer science and networks (192 hours);

Computer Science U.F.R., University of Paris Diderot Fabien de Montgolfier is teaching foundation of computer science, algorithmics, and computer architecture (192 hours);

Master 2 Computer Science, University of Marne-la-Vallée Fabien de Montgolfier is teaching P2P theory and application;

Professional Master, Paris Diderot University Michel Habib is in charge of two courses untitled: Search Engines; Parallelism and mobility, which includes peer-to-peer overlay networks.

6.3. Theses

6.3.1. Defended theses

- Thomas Hugel *Estimations de satisfaisabilité* Jury: Advisors (Directeurs): Y. Boufkhad, M. Habib, Reviewers (Rapporteurs): N. Creignou, L. Kirousis, Examiners (Examineurs): A. Durand, M. Mézard.
- Hoang Anh Phan *Equilibrage de charge et diffusion multicast dans les systèmes pair-à-pair* Jury: Advisor (Directeur): P. Fraigniaud, Reviewers (Rapporteurs): S. Tixeuil, D. Trystram, Examiners (Examineurs): F. Mathieu, L. Viennot.

6.3.2. Ongoing theses

- Hung Tran has begun its PhD since January 2010 *Failure detection with Byzantine adversary*. His advisors are Hugues Fauconnier and Carole Delporte. Hung Tran has a grant from ANR VERSO Shaman.
- Mauricio Soto *Algorithmes de pair à pair et analyse de la structure d'Internet* (Chile-France Allocation).
- Heger Arfaoui *Distributed Computational Complexity*
- Hervé Baumann *protocoles d'échange dans les réseaux distribués*.
- Xavier Koegler *Population protocols*
- Antoine Mamcarz *Algorithmes de décomposition pour les grands graphes* (CNRS grant)

6.3.3. PostDoc

- Julien Clément is in PostDoc (November 2009-November 2011). He works on sensor network and has a grant from Région Ile de France.
- George Giakkoupis works on small world algorithmic aspects and communication complexity.

7. Bibliography

Publications of the year

Articles in International Peer-Reviewed Journal

- [1] C. DELPORTE-GALLET, S. DEVISMES, H. FAUCCONNIER. *Stabilizing leader election in partial synchronous systems with crash failures*, in "J. Parallel Distrib. Comput.", 2010, vol. 70, n^o 1, p. 45-58, <http://dx.doi.org/10.1016/j.jpdc.2009.09.004>.
- [2] C. DELPORTE-GALLET, H. FAUCCONNIER, R. GUERRAOU. *Tight failure detection bounds on atomic object implementations*, in "J. ACM", 2010, vol. 57, n^o 4, <http://doi.acm.org/10.1145/1734213.1734216>.
- [3] Y. EMEK, P. FRAIGNIAUD, A. KORMAN, A. ROSEN. *On the Additive Constant of the k-Server Work Function Algorithm*, in "Information Processing Letters", 2010, vol. 110(24), p. 1120-1123.

- [4] D. FORTIN. *Eigenvectors of Toeplitz matrices under higher order three term recurrence and circulant perturbations*, in "Int. J. Pure Appl. Math.", 2010, vol. 60, n^o 2, p. 217–228.
- [5] D. FORTIN, I. TSEVEENDORJ. *Piece adding technique for convex maximization problems*, in "Journal of Global Optimization", 2010, vol. 48, p. 583-593, 10.1007/s10898-009-9506-z, <http://dx.doi.org/10.1007/s10898-009-9506-z>.
- [6] P. FRAIGNIAUD, A. KORMAN, E. LEBHAR. *Local MST Computation with Short Advice*, in "Theory of Computing Systems", 2010, vol. 47, n^o 14, p. 920-933 [DOI : DOI:10.1007/s00224-010-9280-9], <http://www.ingentaconnect.com/content/klu/224/2010/00000047/00000004/00009280>.
- [7] P. FRAIGNIAUD, E. LEBHAR, Z. LOTKER. *A Lower Bound for Network Navigability*, in "SIAM Journal on Discrete Mathematics", 2010, vol. 24, n^o 1, p. 72-81 [DOI : 10.1137/06067626X], <http://link.ajp.org/link/?SJD/24/72/1>.
- [8] P. FRAIGNIAUD, E. LEBHAR, Z. LOTKER. *Recovering the long-range links in augmented graphs*, in "Theor. Comput. Sci.", 2010, vol. 411, n^o 14-15, p. 1613-1625, <http://dx.doi.org/10.1016/j.tcs.2010.01.009>.
- [9] P. FRAIGNIAUD, D. ILCINKAS, A. PELC. *Communication algorithms with advice*, in "Journal of Computer and System Sciences", 2010, vol. 76, n^o 3-4, See paper for details. [DOI : 10.1016/j.jcss.2009.07.002], http://hal.archives-ouvertes.fr/hal-00412058/PDF/final_JCSS.pdf.
- [10] M. HABIB, C. PAUL. *A survey of the algorithmic aspects of modular decomposition*, in "Computer Science Review", 2010, vol. 4, n^o 1, p. 41-59, <http://dx.doi.org/10.1016/j.cosrev.2010.01.001>.
- [11] A. KORMAN. *Labeling schemes for vertex connectivity*, in "ACM Transactions on Algorithms", 2010, vol. 6, n^o 2, <http://doi.acm.org/10.1145/1721837.1721855>.
- [12] A. KORMAN, D. PELEG, Y. RODEH. *Constructing Labeling Schemes through Universal Matrices*, in "Algorithmica", 2010, vol. 57(4), p. 641-652.

International Peer-Reviewed Conference/Proceedings

- [13] C. DELPORTE-GALLET, S. DEVISMES, H. FAUCONNIER. *Approximation of elta-Timeliness*, in "SSS", S. DOLEV, J. A. COBB, M. J. FISCHER, M. YUNG (editors), Lecture Notes in Computer Science, Springer, 2010, vol. 6366, p. 435-451, http://dx.doi.org/10.1007/978-3-642-16023-3_35.
- [14] C. DELPORTE-GALLET, S. DEVISMES, H. FAUCONNIER, M. LARREA. *Algorithms for Extracting Timeliness Graphs*, in "SIROCCO", B. PATT-SHAMIR, T. EKIM (editors), Lecture Notes in Computer Science, Springer, 2010, vol. 6058, p. 127-141, http://dx.doi.org/10.1007/978-3-642-13284-1_11.
- [15] C. DELPORTE-GALLET, H. FAUCONNIER, R. GUERRAOU, A.-M. KERMARREC. *Brief announcement: byzantine agreement with homonyms*, in "SPAA", F. MEYER AUF DER HEIDE, C. A. PHILLIPS (editors), ACM, 2010, p. 74-75, <http://doi.acm.org/10.1145/1810479.1810494>.
- [16] Y. EMEK, A. KORMAN. *Efficient threshold detection in a distributed environment: extended abstract*, in "29th Annual ACM Symposium on Principles of Distributed Computing (PODC)", Lecture Notes in Computer Science, 2010, p. 183-191.

- [17] P. FRAIGNIAUD, G. GIAKKOUPIS. *On the bit communication complexity of randomized rumor spreading*, in "SPAA", 2010, p. 134-143, <http://doi.acm.org/10.1145/1810479.1810505>.
- [18] P. FRAIGNIAUD, G. GIAKKOUPIS. *On the searchability of small-world networks with arbitrary underlying structure*, in "STOC", 2010, p. 389-398, <http://doi.acm.org/10.1145/1806689.1806744>.
- [19] P. FRAIGNIAUD, A. KORMAN. *An optimal ancestry scheme and small universal posets*, in "42nd ACM Symposium on Theory of Computing (STOC)", Lecture Notes in Computer Science, 2010, p. 611-620.
- [20] P. FRAIGNIAUD, A. KORMAN. *Compact Ancestry Labeling Schemes for XML Trees*, in "21st ACM-SIAM Symposium on Discrete Algorithms (SODA)", 2010, p. 458-466, http://www.siam.org/proceedings/soda/2010/SODA10_038_fraigniaudp.pdf.
- [21] P. FRAIGNIAUD, A. KORMAN. *Compact Ancestry Labeling Schemes for XML Trees*, in "21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)", Lecture Notes in Computer Science, 2010, p. 458-466.
- [22] P. FRAIGNIAUD, A. PELC. *Delays induce an exponential memory gap for rendezvous in trees*, in "SPAA", F. MEYER AUF DER HEIDE, C. A. PHILLIPS (editors), ACM, 2010, p. 224-232, <http://doi.acm.org/10.1145/1810479.1810524>.
- [23] P. FRAIGNIAUD, H. A. PHAN. *"Tree-Farms" for Tree-Based Multicast Schemes in Peer-to-Peer Overlay Networks*, in "2010 IEEE International Conference In Communications (ICC)", 2010.
- [24] P. FRAIGNIAUD, H. A. PHAN. *Degree Hunter: on the Impact of Balancing Node Degrees in de Bruijn-Based Overlay Networks*, in "24th IEEE International Parallel and Distributed Processing Symposium, Workshops and Phd Forum (IPDPSW)", 2010.
- [25] C. GAVOILLE, Q. GODFROY, L. VIENNOT. *Graphes de recouvrement multichemins*, in "12èmes Rencontres Francophones sur les Aspects Algorithmiques de Télécommunications (AlgoTel)", France Belle Dune, Maria Gradinariu Potop-Butucaru and Hervé Rivano, 2010, <http://hal.inria.fr/inria-00475970>.
- [26] C. GAVOILLE, Q. GODFROY, L. VIENNOT. *Multipath Spanners*, in "Structural Information and Communication Complexity, 17th International Colloquium (SIROCCO)", Sirince, 2010, p. 211-223.
- [27] A. KORMAN, S. KUTTEN, D. PELEG. *Proof labeling schemes*, in "Distributed Computing", 2010, vol. 22(4), p. 215-233.
- [28] F. MATHIEU, L. VIENNOT. *Local Aspects of the Global Ranking of Web Pages*, in "10th International Conference on Innovative Internet Community Services (I2CS)", Neuchâtel, 2010, p. 493-506.

Workshops without Proceedings

- [29] Y. BOUFGHAD, T. HUGEL. *Non Uniform Selection of Solutions for Upper Bounding the 3-SAT Threshold*, in "SAT", 2010, p. 99-112.

Scientific Books (or Scientific Book chapters)

- [30] S. DOLEV, J. A. COBB, M. J. FISCHER, M. YUNG (editors). *Stabilization, Safety, and Security of Distributed Systems - 12th International Symposium, SSS 2010, New York, NY, USA, September 20-22, 2010. Proceedings*, Lecture Notes in Computer Science, Springer, 2010, vol. 6366, <http://dx.doi.org/10.1007/978-3-642-16023-3>.
- [31] B. PATT-SHAMIR, T. EKIM (editors). *Structural Information and Communication Complexity, 17th International Colloquium, SIROCCO 2010, Sirince, Turkey, June 7-11, 2010. Proceedings*, Lecture Notes in Computer Science, Springer, 2010, vol. 6058, <http://dx.doi.org/10.1007/978-3-642-13284-1>.
- [32] F. MEYER AUF DER HEIDE, C. A. PHILLIPS (editors). *SPAA 2010: Proceedings of the 22nd Annual ACM Symposium on Parallel Algorithms and Architectures, Thira, Santorini, Greece, June 13-15, 2010*, ACM, 2010.

Other Publications

- [33] Y. BOUFKHAD, T. HUGEL. *Estimating Satisfiability*, 2010, <http://arxiv.org/abs/1005.4518>.
- [34] M. HABIB, J. STACHO. *Unique perfect phylogeny is NP-hard*, 2010, <http://arxiv.org/abs/1011.5737>.