IN PARTNERSHIP WITH:
CNRS

Université Denis Diderot (Paris 7)

Activity Report 2011

Project-Team GANG

Networks, Graphs and Algorithms

IN COLLABORATION WITH: Laboratoire d'Informatique Algorithmique Fondamentale et Appliquée (LIAFA)

RESEARCH CENTER
Paris - Rocquencourt

THEME
Networks and Telecommunications

Table of contents

Project-Team GANG

Keywords: Graph Theory, Dynamic Networks, Fault Tolerance, Peer-to-Peer, Distributed Algorithms

1. Members

Research Scientists

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

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 [Assistant Professor (Paris Diderot Univ.), HdR]
Carole Gallet-Delporte [Professor (Paris Diderot Univ.), HdR]
Fabien de Montgolfier [Assistant professor (Paris Diderot Univ.)]

PhD Students

Heger Arfaoui [PhD Student]
Hervé Baumann [PhD Student]
Xavier Koegler [PhD Student]
Antoine Mamcarz [PhD Student]
Mauricio Soto [PhD Student]
Thu-Hien To [PhD Student]
Hung Tran-The [PhD Student]

Administrative Assistant

Christine Anocq [TRE shared time 20% (with Hipercom and Aoste)]

2. Overall Objectives

2.1. Introduction

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 theoretically 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.

2.2. Highlights

The paper [22] was awarded as a best article in the *25th Int. Symp. on Distributed Computing (DISC 2011)*.

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. Distributed algorithms without knowledge of global parameters

Participants: Amos Korman, Jean-Sébastien Sereni, Laurent Viennot.

Many fundamental local distributed algorithms are non-uniform, that is, they assume that all nodes know good estimations of one or more global parameters of the network, e.g., the maximum degree Δ or the number of nodes n. In [28], we introduce a rather general technique for transforming a non-uniform algorithm into a uniform one with same asymptotic complexity.

4.1.2. Asymptotic modularity

Participants: Fabien de Montgolfier, Mauricio Soto, Laurent Viennot.

Modularity has been introduced as a quality measure for graph partitioning by Newman and Girvan. It has received considerable attention in several disciplines, especially complex systems. In order to better understand this measure from a graph theoretical point of view, we study in [32], [31] the asymptotic modularity of a variety of graph classes.

4.1.3. Internet Structure

Participants: Fabien de Montgolfier, Mauricio Soto, Laurent Viennot.

In [33], [1], we study the measurement of the Internet according to two graph parameters: treewidth and hyperbolicity.

4.1.4. Multipath Spanners

Participants: Cyril Gavoille, Quentin Godfroy, Laurent Viennot.

Motivated by multipath routing, we introduce in [23], [39] a multi-connected variant of spanners.

4.1.5. $\delta-$hyperbolicity

Participants: Victor Chepoi [CNRS LIF, University of Marseille, France], Feodor Dragan [University of Ohio, USA], Bernard Estrellon [CNRS LIF, University of Marseille, France], Michel Habib [CNRS LIAFA, University of Paris Diderot, France], Yann Vaxes [University of Florence, Italy], Yang Xiang [University of Ohio, USA].

$\delta-$Hyperbolic metric spaces have been defined by M. Gromov in 1987 via a simple 4-point condition: for any four points u, v, w, x, the two larger of the distance sums $d(u,v)+d(w,x), d(u,w)+d(v,x), d(u,x)+d(v,w)$ differ by at most 2δ. They play an important role in geometric group theory, geometry of negatively curved spaces, and have recently become of interest in several domains of computer science, including algorithms and networking. In [5] paper, we study unweighted $\delta-$hyperbolic graphs. Using the Layering Partition technique, we show that every $n-$vertex δ-hyperbolic graph with $\delta \geq 1/2$ has an additive $O(\delta \log n)-$spanner with at most $O(\delta n)$ edges and provide a simpler, in our opinion, and faster con- struction of distance approximating trees of δ-hyperbolic graphs with an additive error $O(\delta \log n)$. The construction of our tree takes only linear time in the size of the input graph. As a consequence, we show that the family of $n-$vertex $\delta-$hyperbolic graphs with $\delta \geq 1/2$ admits a routing labeling scheme with $O(\delta \log^2 n)$ bit labels, $O(\delta \log n)$ additive stretch and $O(\log_2(4\delta))$ time routing protocol, and a distance labeling scheme with $O(\log^2 n)$ bit labels, $O(\delta \log n)$ additive error and constant time distance decoder.

4.1.6. Perfect Phylogeny

4.1.6.1. Perfect Phylogeny Is $NP-$ Hard

Participants: Michel Habib [CNRS LIAFA, University of Paris Diderot, France], Juraj Stacho [University of Haifa, Israel].

We answer in the affirmative [24], to the question pro- posed by Mike Steel as a $100 challenge: "Is the following problem $NP-$ hard? Given a ternary 1 phylogenetic X-tree T and a collection Q of quartet subtrees on X, is T the only tree that displays Q?" As a particular consequence of this, we show that the unique chordal sandwich problem is also $NP-$hard.

4.1.6.2. Compatibility of Multi-states Characters

Participants: Michel Habib [CNRS LIAFA, University of Paris Diderot, France], Thu-Hien To [CNRS LIAFA, University of Paris Diderot, France].

Perfect phylogeny consisting of determining the compatibil- ity of a set of characters is known to be $NP-$complete. We propose in [25], a conjecture on the necessary and sufficient conditions of compatibility: Given a set C of $r-$states full characters, there exists a function $f(r)$ such that C is compatible iff every set of $f(r)$ characters of C is compatible. According to numerous references, $f(2)=2$, $f(3)=3$ and $f(r) \geq r$. Some conjectured that $f(r)=r$ for any $r \geq 2$. In this paper, we present an example showing that $f(4) \geq 5$. Therefore it could be the case that for $r \geq 4$ characters, the problem behavior drastically changes. In a second part, we propose a closure operation for chordal sandwich graphs. The later problem is a common approach of perfect phylogeny.

4.1.7. Graph sandwich

Participants: Arnaud Durand [CNRS LIAFA, University of Paris Diderot, France], Michel Habib [CNRS LIAFA, University of Paris Diderot, France].

Graph sandwich problems were introduced by Golumbic et al. (1994) in [12] for DNA physical mapping problems and can be described as follows. Given a property Π of graphs and two disjoint sets of edges E_1 , E_2 with $E_1 \subseteq E_2$ on a vertex set V , the problem is to find a graph G on V with edge set E_s having property Π and such that $E_1 \subseteq E_s \subseteq E_2$. In [8] paper, we exhibit a quasi-linear reduction between the problem of finding an independent set of size $k \geq 2$ in a graph and the problem of finding a sandwich homogeneous set of the same size k. Using this reduction, we prove that a number of natural (decision and counting) problems related to sandwich homogeneous sets are hard in general. We then exploit a little further the reduction and show that finding efficient algorithms to compute small sandwich homogeneous sets would imply substantial improvement for computing triangles in graphs.

4.1.8. Diameter of Real-World Undirected Graphs

Participants: Pierluigi Crescenzi [University of Florence, Italy], Roberto Grossi [University of Pisa, Italy], Michel Habib [CNRS LIAFA, University of Paris Diderot, France], Lorenzo Lanzi [University of Florence, Italy], Andrea Marino [University of Florence, Italy].

In [16], we propose a new algorithm for computing the diameter of undirected unweighted graphs. Even though, in the worst case, this algorithm has complexity $O(nm)$, where n is the number of nodes and m is the number of edges of the graph, we experimentally show (on almost 200 real-world graphs) that in practice our method works in linear time. Moreover, we show how to extend our algorithm to the case of undirected weighted graphs and, even in this case, we present some preliminary very positive experimental results.

4.1.9. Parsimonious flooding in dynamic graphs

Participants: Hervé Baumann [CNRS LIAFA, University of Paris Diderot, France], Pierluigi Crescenzi [University of Florence, Italy], Pierre Fraigniaud [CNRS LIAFA, University of Paris Diderot, France].

An edge-Markovian process with birth-rate p and death-rate q generates infinite sequences of graphs $(G_0, G_1, G_2, \ldots)$ with the same node set $[n]$ such that G_t is obtained from G_{t-1} as follows: if $e \notin E(G_{t-1})$ then $e \in E(G_t)$ with probability p, and if $e \in E(G_{t-1})$ then $e \notin E(G_t)$ with probability q. In [2], we establish tight bounds on the complexity of flooding in edge-Markovian graphs, where flooding is the basic mechanism in which every node becoming aware of an information at step t forwards this information to all its neighbors at all forthcoming steps $t' > t$. These bounds complete previous results obtained by Clementi et al. Moreover, we also show that flooding in dynamic graphs can be implemented in a parsimonious manner, so that to save bandwidth, yet preserving efficiency in term of simplicity and completion time. For a positive integer k, we say that the flooding protocol is $k-$active if each node forwards an information only during the k time steps immediately following the step at which the node receives that information for the first time. We define the reachability threshold for the flooding protocol as the smallest integer k such that, for any source s[n] , the $k-$active flooding protocol from s completes (i.e., reaches all nodes), and we establish tight bounds for this parameter. We show that, for a large spectrum of parameters p and q, the reachability threshold is by several orders of magnitude smaller than the flooding time. In particular, we show that it is even constant whenever the ratio $p/(p+q)$ exceeds $\log n/n$. Moreover, we also show that being active for a number of steps equal to the reachability threshold (up to a multiplicative constant) allows the flooding protocol to complete in optimal time, i.e., in asymptotically the same number of steps as when being perpetually active. These results demonstrate that flooding can be implemented in a practical and efficient manner in dynamic graphs. The main ingredient in the proofs of our results is a reduction lemma enabling to overcome the time dependencies in edge-Markovian dynamic graphs.

4.2. Distributed computational complexities

4.2.1. Local Distributed Decision

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

Inspired by sequential complexity theory, in [20] we focus on a complexity theory for distributed decision problems. We first study the intriguing question of whether randomization helps in local distributed computing, and to what extent. Our main result provides a sharp threshold for the impact of randomization on decision hereditary problems. In addition, we investigate the impact of non-determinism on local decision, and establish some structural results inspired by classical computational complexity theory. Specifically, we show that non-determinism does help, but that this help is limited, as there exist languages that cannot be decided non-deterministically. Perhaps surprisingly, it turns out that it is the combination of randomization with non-determinism that enables to decide all languages in constant time. Finally, we introduce the notion of local reduction, and establish some completeness results

4.2.2. Asynchronous Wait-free Decision

Participants: Pierre Fraigniaud [CNRS LIAFA, University of Paris Diderot, France], Sergio Rajsbaum [Maths. Institute, University of Mexico, Mexico], Corentin Travers [Technion, Israel].

In order to capture the core of asynchronous distributed decision model, we address in [22] the *wait-free* model with crash failures. The set of tasks whose input is a pair (s, t) and deciding whether $t \in \Delta(s)$, i.e. whether t is a valid output for s, has been proven to be decidable in this model.

4.2.3. Mobile Distributed Decision

Participants: Pierre Fraigniaud [CNRS LIAFA, University of Paris Diderot, France], Andrzej Pelc [UQO, University of Quebec, Canada].

In [21], we partially answer the question of decidability of any language for mobile agents in a 2D environment like telecom networks or robots. It is proven that, for every agent, verifying whether (i) he/she is alone or not and (ii) he/she is able to capture the environment, is associated with the question of pertaining to an equivalence class of a map. A positive answer helps in the non-deterministic decision for any language for mobile agent.

4.2.4. Approximating the Statistics of various Properties in Randomly Weighted Graphs

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

In [19], we consider the setting of randomly weighted graphs. Under this setting, weighted graph properties typically become random variables and we are interested in computing their statistical features. Unfortunately, this turns out to be computationally hard for some weighted graph properties albeit the problem of computing the properties per se in the traditional setting of algorithmic graph theory is tractable. For example, there are well known efficient algorithms that compute the diameter of a given weighted graph, yet, computing the expected diameter of a given randomly weighted graph is ♯P-hard even if the edge weights are identically distributed. In this paper, we define a family of weighted graph properties and show that for each property in this family, the problem of computing the k'th moment (and in particular, the expected value) of the corresponding random variable in a given randomly weighted graph G admits a fully polynomial time randomized approximation scheme (FPRAS) for every fixed k. This family includes fundamental weighted graph properties such as the diameter of G, the radius of G (with respect to any designated vertex) and the weight of a minimum spanning tree of G.

4.2.5. New bounds for the controller problem

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

In [10], we establish two new lower bounds on the message complexity of the controller problem. We first prove a simple lower bound stating that any (M, W)-controller must send $\Omega(N \log \frac{M}{W+1})$ messages. Second, for the important case when W is proportional to M (this is the common case in most applications), we use a surprising reduction from the (centralized) monotonic labeling problem to show that any (M, W)-controller must send $\Omega(N \log N)$ messages. In fact, under a long lasting conjecture regarding the complexity of the monotonic labeling problem, this lower bound is improved to a tight $\Omega(N \log^2 N)$.

4.2.6. Online computation with advice

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 Rosén [CNRS LIAFA, University of Paris Diderot, France].

In [9], we consider a model for online computation in which the online algorithm receives, together with each request, some information regarding the future, referred to as advice. We are interested in the impact of such advice on the competitive ratio, and in particular, in the relation between the size b of the advice, measured in terms of bits of information per request, and the (improved) competitive ratio. In this paper we propose the above model and illustrate its applicability by considering two of the most extensively studied online problems, namely, metrical task systems (MTS) and the k-server problem.

4.2.7. Tight Bounds For Distributed MST Verification

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

In [26], we establishes tight bounds for the Minimum-weight Spanning Tree (MST) verification problem in the distributed setting. Specifically, we provide an MST verification algorithm that achieves simultaneously $\widetilde{O}(|E|)$ messages and $\widetilde{O}(\sqrt{n}+D)$ time, where $|E|$ is the number of edges in the given graph G and D is G's diameter. On the negative side, we show that any MST verification algorithm must send $\Omega(|E|)$ messages and incur $\widetilde{\Omega}(\sqrt{n}+D)$ time in worst case. Our upper bound result appears to indicate that the verification of an MST may be easier than its construction.

4.2.8. Distributed verification and hardness of distributed approximation

Participants: Atish Das Sarma [Google research, USA], Stephan Holzer [ETH, Zurich, Switzerland], Liah Kor [Weizmann Institute of Science, Israel], Amos Korman [CNRS LIAFA, University of Paris Diderot, France], Danupon Nanongkai [Nanyang Technological University, Singapore], David Peleg [Weizmann Institute of Science, Israel], Roger Wattenhofer [ETH, Zurich, Switzerland].

In [30], we initiate a systematic study of distributed verification, and give almost tight lower bounds on the running time of distributed verification algorithms for many fundamental problems such as connectivity, spanning connected subgraph, and $s-t$ cut verification. We then show applications of these results in deriving strong unconditional time lower bounds on the hardness of distributed approximation for many classical optimization problems including minimum spanning tree, shortest paths, and minimum cut. Many of these results are the first non-trivial lower bounds for both exact and approximate distributed computation and they resolve previous open questions. Moreover, our unconditional lower bound of approximating minimum spanning tree (MST) subsumes and improves upon the previous hardness of approximation bound of Elkin [STOC 2004] as well as the lower bound for (exact) MST computation of Peleg and Rubinovich [FOCS 1999]. Our result implies that there can be no distributed approximation algorithm for MST that is significantly faster than the current exact algorithm, for any approximation factor. Our lower bound proofs show an interesting connection between communication complexity and distributed computing which turns out to be useful in establishing the time complexity of exact and approximate distributed computation of many problems.

4.3. Peer to Peer Networks Performance

Participants: Fabien Mathieu, François Baccelli.

In [3], we present and discuss possible architectures for P2P systems to manage overlays that try to cope with the underlying network.

In [40], [29], we discuss theoretical performance issues that arise from using “Live Seeding”, a technique that can be employed to leverage the capacity of a P2P/Hybrid Live Streaming Systems by utilizing the capacities of idle peers.

In [38], we propose a new paradigm for P2P networks, where the bandwidth bottleneck is not the access node anymore. This new model is versatile enough to be used in the context of classical networks with congestion control, wireless networks, or semantic networks.

4.4. Fault Tolerance in Distributed Networks

4.4.1. Verification of population protocols

Participants: Hugues Fauconnier, Carole Gallet-Delporte.

In [15], we address the problem of verification by model- checking of the basic population protocol (PP) model of Angluin et al. . This problem has received special attention in the last two years and new tools have been proposed to deal with it. We show that the problem can be solved by using the existing model-checking tools, e.g., Spin and Prism. In order to do so, we apply the counter abstraction to get an abstraction of the PP model which can be efficiently verified by the existing model-checking tools. Moreover, this abstraction preserves the correct stabilization property of PP models. To deal with the fairness assumed by the PP models, we provide two new recipes. The first one gives sufficient conditions under which the PP model fairness can be replaced by the weak fairness implemented in Spin. We show that this recipe can be applied to several PP models. In the second recipe, we show how to use probabilistic model-checking and, in particular, Prism to take completely in consideration the fairness of the PP models. The correctness of this recipe is based on existing theorems involving finite discrete Markov chains. An abstract of this paper has been also published in [34].

4.4.2. Failure Detection

Participants: Hugues Fauconnier, Carole Gallet-Delporte.

What does it mean to solve a distributed task? In Paxos, Lamport proposed a definition of solvability in which every process is split into a proposer that submits commands to be executed, an acceptor that takes care of the command execution order, and a learner that receives the outcomes of executed commands. The resulting perspective of computation in which every proposed command can be executed, be its proposer correct or faulty, proved to be very useful when processes take steps on behalf of each other, i.e., in *simulations*.

Most interesting tasks cannot be solved asynchronously, and failure detectors were proposed to circumvent these impossibilities. Alas, when it comes to solving a task using a failure detector, we cannot leverage simulation-based techniques. A process cannot perform steps of failure detector-based computation on behalf of another process, since it cannot access the remote failure-detector module.

In [17], we propose a new definition of solving a task with a failure detector in which *computation* processes that propose inputs and provide outputs are treated separately from *synchronization* processes that coordinate using a failure detector. In the resulting framework, any failure detector is shown to be equivalent to the availability of some k-set agreement. As a corollary, we obtain a complete classification of tasks, including ones that evaded comprehensible characterization so far, such as renaming.

Shared objects like atomic register, test-and-set, cmp-and-swap are classical hardware primitives that help to develop fault-tolerant distributed applications. In order to compare shared objects, in [41], we consider their implementations in message passing models. With the minimal failure detector for each object, we get a new hierarchy that has only two levels. This paper summarizes recent works and results on this topic.

In [7], we first define the basic notions of *local* and *non-local* tasks for distributed systems. Intuitively, a task is local if, in a system with no failures, each process can compute its output value locally by applying some local function on its own input value (so the output value of each process depends only on the process' own input value, not on the input values of the other processes); a task is non-local otherwise. All the interesting distributed tasks, including all those that have been investigated in the literature (e.g., consensus, set agreement, renaming, atomic commit, etc.) are non-local.

In this paper we consider non-local tasks and determine the minimum information about failures that is necessary to solve such tasks in message-passing distributed systems. As part of this work, we also introduces *weak set agreement* — a natural weakening of *set agreement* — and show that, in some precise sense, it is the weakest non-local task in message-passing systems.

4.4.3. Adversary disagreement and Byzantine agreement

Participants: Hugues Fauconnier, Carole Gallet-Delporte.

At the heart of distributed computing lies the fundamental result that the level of agreement that can be obtained in an asynchronous shared memory model where t processes can crash is exactly t + 1. In other words, an adversary that can crash any subset of size at most t can prevent the processes from agreeing on t values. But what about all the other $2^{2^n-1} - (n+1)$ adversaries that are not uniform in this sense and might crash certain combination of processes and not others? In [6], we present a precise way to classify all adversaries. We introduce the notion of disagreement power: the biggest integer k for which the adversary can prevent processes from agreeing on k values. We show how to compute the disagreement power of an adversary and derive n equivalence classes of adversaries.

So far, the distributed computing community has either assumed that all the processes of a distributed system have distinct identifiers or, more rarely, that the processes are anonymous and have no identifiers. These are two extremes of the same general model: namely, n processes use ℓ different authenticated identifiers, where $1 \leq \ell \leq n$. In [18], we ask how many identifiers are actually needed to reach agreement in a distributed system with t Byzantine processes.

We show that having $3t+1$ identifiers is necessary and sufficient for agreement in the synchronous case but, more surprisingly, the number of identifiers must be greater than $\frac{n+3t}{2}$ in the partially synchronous case. This demonstrates two differences from the classical model (which has $\ell = n$): there are situations where relaxing synchrony to partial synchrony renders agreement impossible; and, in the partially synchronous case, increasing the number of *correct* processes can actually make it harder to reach agreement. The impossibility proofs use the fact that a Byzantine process can send multiple messages to the same recipient in a round. We show that removing this ability makes agreement easier: then, $t+1$ identifiers are sufficient for agreement, even in the partially synchronous model.

4.4.4. Fast and compact self stabilizing verification, computation, and fault detection of an MST

Participants: Amos Korman [CNRS LIAFA, University of Paris Diderot, France], Shay Kutten [Technion, Israel], Toshimitsu Masuzawa [Osaka University, Japan].

In [27], we address the impact of optimizing the memory size on the time complexity, and show that this carries at most a small cost in terms of time in the context of MST. Specifically, we present a self stabilizing distributed verification algorithm whose time complexity is $O(\log^2 n)$ in synchronous networks, or $O(\Delta \log^2 n)$ in asynchronous networks, where Δ denotes the largest degree of a node. More importantly, the memory size at each node remains optimal- $O(\log n)$ bits throughout the execution. This answers an open problem posed by Awerbuch and Varghese (FOCS 1991). We also show that $\Omega(\log n)$ time is necessary if the memory size is restricted to $O(\log n)$ bits, even in synchronous networks. We demonstrate the usefulness of our verification scheme by using it as a module in a new self stabilizing MST construction algorithm. This algorithm has the important property that, if faults occur after the construction ended, they are detected by some nodes within $O(\log^2 n)$ time in synchronous networks, or within $O(\Delta \log^2 n)$ time in asynchronous networks. The rest of the nodes detect within $O(D \log n)$ time, where D denotes the diameter. Moreover, if a constant number of faults occur, then, within the required detection time above, they are detected by some node in the $O(\log n)$ locality of each of the faults. The memory size of the self stabilizing MST construction is $O(\log n)$ bits per node (optimal), and the time complexity is $O(n)$. This time complexity is significantly better than the best time complexity of previous self stabilizing MST algorithms, that was $\Omega(n^2)$ even when using memory of $\Omega(\log^2 n)$ bits, and even without having the above localized fault detection property.The time complexity of previous algorithms that used $O(\log n)$ memory size was $O(n|E|)$.

4.5. Discrete Optimization Algorithms

4.5.1. Estimating Satisfiability

Participants: Yacine Boufkhad, Thomas Hugel.

In [4], the problem of estimating the proportion of satisfiable instances of a given CSP (constraint satisfaction problem) can be tackled through weighting. It 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. We define in this paper a general weighting scheme for the estimation of satisfiability of general CSPs. First we give some sufficient conditions for a weighting system to be correct. Then we show that this scheme allows for an improvement on the upper bound on the existence of non-trivial cores in 3-SAT obtained by Maneva and Sinclair (2008) to 4.419. Another more common way of estimating satisfiability is ordering. This 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. We compare ordering and weighting under various conditions.

4.5.2. Eigenvectors of three term recurrence Toeplitz matrices and Riordan group

Participant: Dominique Fortin.

Eigenvalues of tridiagonal (including main) Toeplitz matrices are analytically known under some regular distance to the main diagonal. Any eigenvector may be easily computed then, through a backward process; instead, in [11], we give an analytical form for each component through the reciprocation of the underlied trinomial. More generally, the connection to the Riordan group follows some bilinear iterative process.

4.5.3. Piecewise Convex Maximization problems and algorithms

Participants: Dominique Fortin, Ider Tseveendorj.

In [14], we provide a global search algorithm for maximizing a piecewise convex function F over a compact D. We propose to iteratively refine the function F at local solution y by a *virtual cutting* function $p_y(\cdot)$ and to solve
$\max\{\min\{F(x) - F(y), p_y(x)\} \mid x \in D\}$ instead. We call this function either a patch, when it avoids returning back to the same local solutions, or a pseudo patch, when it possibly yields a better point. It is *virtual* in the sense that the role of cutting constraints is played by additional convex pieces in the objective function. We report some computational results, that represent an improvement on previous linearization based techniques.

It is well known that maximization of any difference of convex functions could be turned into a convex maximization; in [13], we aim at a piecewise convex maximization problem instead. Despite, it may seem harder, sometimes the dimension may be reduced by 1 and the local search improved by using extreme points of the closure of the convex hull of better points. We show that it is always the case for both binary and permutation problems and give, as such instances, piecewise convex formulations for the maximum clique problem and the quadratic assignment problem.

in [12], we consider mathematical programming problems with the so-called piecewise convex objective functions. A solution method for this interesting and important class of nonconvex problems is presented. This method is based on Newton's law of universal gravitation, multicriteria optimization and Helly's theorem on convex bodies. Numerical experiments using well known classes of test problems on piecewise convex maximization, convex maximization as well as the maximum clique problem show the efficiency of the approach.

5. Contracts and Grants with Industry

5.1. Contracts with Industry

Collaboration with Alcatel-Lucent Bell Labs France (ALBLF)

Within the Laboratory of Information, Networking and Communication Sciences (LINCS), collaborations have been made with ALBLF. In 2011, it resulted in two internships paid by ALBLF and co-supervised by Fabien Mathieu (INRIA) and Ludovic Noirie (ALBLF). In 2012, both interns should start a thesis in collaboration with ALBLF and INRIA (one CIFRE, one in the context of the joint lab).

6. Partnerships and Cooperations

6.1. Regional Initiatives

6.1.1. 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.

6.2. National Initiatives

6.2.1. ANR Algorithm Design and Analysis for Implicitly and Incompletely Defined Interaction Networks (ALADDIN)

Participants: Cyril Gavoille [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). Ending in november 2011, the project is prolonged until end of 2012 for LABRI partner.

6.2.2. ANR PROSE

Participants: Pierre Fraigniaud, Amos Korman, Laurent Viennot.

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

6.2.3. ANR Shaman

Participants: Hugues Fauconnier, Pierre Fraigniaud, Carole Gallet-Delporte, Hung Tran-The, Laurent Viennot.

Managed by University Paris Diderot, H. Fauconnier leads this project that grants Ph. D. H. Tran-The.

6.2.4. ANR Displexity

Participants: Hugues Fauconnier, Pierre Fraigniaud, Carole Gallet-Delporte, Amos Korman, Hung Tran-The, Laurent Viennot.

Managed by University Paris Diderot, C. Delporte and H. Fauconnier lead this project that grants 1 Ph. D. and 2 internships per year.

6.3. European Initiatives

6.3.1. FP7 Projet

6.3.1.1. EULER

Title: 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)

See also: http://www.euler-fire-project.eu/

Abstract: 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.

6.3.2. Collaborations in European Programs, except FP7

Program: EIT ICT Labs

Project acronym: TREC-EIT-GA2011-HORS-5643

Project title:

Duration: 2011

Coordinator:Ilkka Norros

Other partners: KTH (Finland), Fraunhofer (Germany)

Abstract: Content Distribution challenging issues; managed by TREC for France, the project allowed Pascal Felber to be invited by Fabien Mathieu for a postdoctoral position.

6.4. Teaching

Master MPRI

- Michel Habib is in charge of a course entitled “graph algorithms”.
- Pierre Fraigniaud (12 hours) is in charge of the course “Algorithmique distribuée pour les réseaux”;
- Carole Delporte and Hugues Fauconnier are in charge of “Algorithmique distribuée avec mémoire partagée”;
- Laurent Viennot (12 hours) is teaching “Structures de données distribuées et routage”

D.U.T., University of Paris Diderot

- Yacine Boufkhad (192 hours) is teaching scientific computer science and networks.

Computer Science U.F.R., University of Paris Diderot

- Fabien de Montgolfier (192 hours) 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 (192 hours) is in charge of two courses entitled: Search Engines; Parallelism and mobility, which includes peer-to-peer overlay networks;
- Carole Delporte (192 hours) is teaching “Distributed programming”;
- Hugues Fauconnier (192 hours) in charge of both courses “Internet Protocols and Distributed algorithms”.

Master 2 Computer Science, University of Paris 6

- Fabien Mathieu is teaching Peer-to-peer Networks (6 hours).

PhD : Mauricio Soto, ”Quelques propriétés topologiques des graphes et applications à Internet et aux réseaux”, Paris Diderot University, 2 December 2011, supervisors: Fabien de Montgolfier et Laurent Viennot;

PhD : Thu-Hien To: ”On some graph problems in phylogenetics”, Paris Diderot University, 15 September 2011, supervisor: Michel Habib;

PhD in progress : Hung Tran-The, Failure detection with Byzantine adversary, from 2010, supervisors: Hugues Fauconnier and Carole Delporte,

7. Bibliography

Publications of the year

Doctoral Dissertations and Habilitation Theses

[1] M. SOTO. *Quelques propriétés topologiques des graphes et applications à Internet et aux réseaux*, Paris Diderot, Dec 2011.

Articles in International Peer-Reviewed Journal

[2] H. BAUMANN, P. CRESCENZI, P. FRAIGNIAUD. *Parsimonious flooding in dynamic graphs*, in "Distributed Computing", 2011, vol. 24, n^{o} 1, p. 31-44.

[3] R. BIRKE, E. LEONARDI, M. MELLIA, A. BAKAY, T. SZEMETHY, C. KIRALY, R. LO CIGNO, F. MATHIEU, L. MUSCARIELLO, S. NICCOLINI, J. SEEDORF, G. TROPEA. *Architecture of a Network-Aware P2P-TV Application: The NAPA-WINE Approach*, in "IEEE Communications Magazine", 2011, vol. 49, n^{o} 6, p. 154-163.

[4] Y. BOUFKHAD, T. HUGEL. *Estimating satisfiability*, in "Discrete Applied Mathematics", 2012, vol. 160, n^{o} 1-2, p. 61 - 80, à paraitre [*DOI :* 10.1016/J.DAM.2011.10.005], http://www.sciencedirect.com/science/article/pii/S0166218X11003647.

[5] V. CHEPOI, F. DRAGAN, B. ESTRELLON, M. HABIB, Y. VAXES, Y. XIANG. *Additive Spanners and Distance and Routing Labeling Schemes for Hyperbolic Graphs*, in "Algorithmica", 2011, vol. on line first.

[6] C. DELPORTE-GALLET, H. FAUCONNIER, R. GUERRAOUI, A. TIELMANN. *The disagreement power of an adversary*, in "Distributed Computing", 2011, vol. 24, n^{o} 3-4, p. 137-147.

[7] C. DELPORTE-GALLET, H. FAUCONNIER, S. TOUEG. *The Minimum Information about Failures for Solving non-local Tasks in Message-Passing Systems*, in "Distributed Computing", 2011, vol. On line first.

[8] A. DURAND, M. HABIB. *Complexity issues for the sandwich homogeneous set problem*, in "Discrete Applied Mathematics", 2011, vol. 159, n^{o} 7, p. 574-580.

[9] Y. EMEK, P. FRAIGNIAUD, A. KORMAN, A. ROSÉN. *Online computation with advice*, in "Theor. Comput. Sci", 2011, vol. 412, n^{o} 24, p. 2642-2656.

[10] Y. EMEK, A. KORMAN. *New bounds for the controller problem*, in "Journal of Distributed Computing", 2011, vol. 24, n^{o} 3-4, p. 177-186.

[11] D. FORTIN. *Eigenvectors of three term recurrence Toeplitz matrices and Riordan group*, in "Int. J. Pure Appl. Math.", 2011, vol. 67, n^{o} 1, p. 97–105.

[12] D. FORTIN, I. TSEVEENDORJ. *Attractive force search algorithm for piecewise convex maximization problems*, in "Opt. lett. Online First", 15 September 2011, n^{o} 145, http://www.springerlink.com/content/1862-4472/preprint/?sort=p_OnlineDate&sortorder=asc&o=140.

[13] D. FORTIN, I. TSEVEENDORJ. *Piecewise Convex formulations of binary and permutation problem*, in "Siberian Journal of Numerical Mathematics", 2011, vol. 14, n^{o} 4, p. 409–423.

[14] D. FORTIN, I. TSEVEENDORJ. *Piecewise convex maximization problems: piece adding technique*, in "J. Optim. Theory Appl.", 2011, vol. 148, n^{o} 3, p. 471–487, http://dx.doi.org/10.1007/s10957-010-9763-5.

International Conferences with Proceedings

[15] J. CLÉMENT, C. DELPORTE-GALLET, H. FAUCONNIER, M. SIGHIREANU. *Guidelines for the Verification of Population Protocols*, in "ICDCS", IEEE Computer Society, 2011, p. 215-224.

[16] P. CRESCENZI, R. GROSSI, M. HABIB, L. LANZI, A. MARINO. *On Computing the Diameter of Real-World Undirected Graphs*, in "Workshop ESA in honor of Giorgio Ausiello, submitted to TCS", 2011.

[17] C. DELPORTE-GALLET, H. FAUCONNIER, E. GAFNI, P. KUZNETSOV. *Brief Announcement: On the Meaning of Solving a Task with a Failure Detector*, in "DISC", D. PELEG (editor), Lecture Notes in Computer Science, Springer, 2011, vol. 6950, p. 145-146.

[18] C. DELPORTE-GALLET, H. FAUCONNIER, R. GUERRAOUI, A.-M. KERMARREC, E. RUPPERT, H. TRAN-THE. *Byzantine agreement with homonyms*, in "PODC", San Jose, CA, USA, C. GAVOILLE, P. FRAIGNIAUD (editors), ACM, June 2011, p. 21-30.

[19] Y. EMEK, A. KORMAN, Y. SHAVITT. *Approximating the Statistics of various Properties in Randomly Weighted Graphs*, in "22nd Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA)", San Francisco, 2011, p. 1455-1467.

[20] P. FRAIGNIAUD, A. KORMAN, D. PELEG. *Local Distributed Decision*, in "52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)", Palm Springs, 2011.

[21] P. FRAIGNIAUD, A. PELC. *Delays Induce an Exponential Memory Gap for Rendezvous in Trees*, in "10th Latin American Theoretical Informatics Symposium (LATIN)", Springer, 2012, à paraitre.

[22] P. FRAIGNIAUD, S. RAJSBAUM, C. TRAVERS. *Locality and Checkability in Wait-Free Computing*, in "25th International Symposium on Distributed Computing (DISC)", LNCS, Springer, 2011, vol. 6950, p. 333-347.

[23] C. GAVOILLE, Q. GODFROY, L. VIENNOT. *Node-Disjoint Multipath Spanners and their Relationship with Fault-Tolerant Spanners*, in "15th International Conference on Principles of Distributed Systems (OPODIS)", Toulouse, 2011, p. 1-16.

[24] M. HABIB, J. STACHO. *Unique Perfect Phylogeny Is NP-Hard*, in "CPM", 2011, p. 132-146.

[25] M. HABIB, T.-H. TO. *On a Conjecture about Compatibility of Multi-states Characters*, in "WABI", 2011, p. 116-127.

[26] L. KOR, A. KORMAN, D. PELEG. *Tight Bounds For Distributed MST Verification*, in "28th International Symposium on Theoretical Aspects of Computer Science, (STACS)", Dortmund, 2011, p. 69-80.

[27] A. KORMAN, S. KUTTEN, T. MASUZAWA. *Fast and compact self stabilizing verification, computation, and fault detection of an MST*, in "30th Annual ACM Symposium on Principles of Distributed Computing (PODC)", San Jose, ACM Press, 2011, p. 311-320.

[28] A. KORMAN, J.-S. SERENI, L. VIENNOT. *Toward more Localized Local Algorithms: Removing Assumptions concerning Global Knowledge*, in "30th Annual ACM Symposium on Principles of Distributed Computing (PODC)", San Jose, 2011, p. 49-58.

[29] F. MATHIEU. *Live Seeding: Performance Bounds of Seeders for P2P Live Streaming*, in "P2P '11: Proceedings of the 2011 Eleventh International Conference on Peer-to-Peer Computing", IEEE Computer Society, 2011, p. 172–181.

[30] A. D. SARMA, S. HOLZER, L. KOR, A. KORMAN, D. NANONGKAI, G. PANDURANGAN, D. PELEG, R. WATTENHOFER. *Distributed verification and hardness of distributed approximation*, in "43rd ACM Symposium on Theory of Computing, (STOC)", San Jose, ACM Press, 2011, p. 363-372.

[31] F. DE MONTGOLFIER, M. SOTO, L. VIENNOT. *Asymptotic Modularity of some Graph Classes*, in "22nd International Symposium on Algorithms and Computation (ISAAC)", Yokohama, 2011, p. 435–444.

[32] F. DE MONTGOLFIER, M. SOTO, L. VIENNOT. *Clustering de métrique et clustering de graphe*, in "13es Rencontres Francophones sur les Aspects Algorithmiques de Télécommunications (AlgoTel)", Cap Estérel, France, P. DUCOURTHIAL (editor), 2011, Equipe-projet GANG (inria), http://hal.inria.fr/inria-00583844/en/.

[33] F. DE MONTGOLFIER, M. SOTO, L. VIENNOT. *Treewidth and Hyperbolicity of the Internet*, in "10th IEEE International Symposium on Network Computing and Applications (IEEE NCA)", Boston, 2011, p. 1–8.

National Conferences with Proceeding

[34] J. CLÉMENT, C. DELPORTE-GALLET, H. FAUCONNIER, M. SIGHIREANU. *Mode d'emploi pour la vérification des protocoles de population*, in "13es Rencontres Francophones sur les Aspects Algorithmiques de Télécommunications (AlgoTel)", Cap Estérel, France, P. DUCOURTHIAL (editor), 2011, http://hal.inria.fr/inria-00584684/en/.

Scientific Books (or Scientific Book chapters)

[35] *2011 International Conference on Distributed Computing Systems, ICDCS 2011, Minneapolis, Minnesota, USA, June 20-24, 2011*, IEEE Computer Society, 2011.

[36] D. PELEG (editor). *Distributed Computing - 25th International Symposium, DISC 2011, Rome, Italy, September 20-22, 2011. Proceedings*, Lecture Notes in Computer Science, Springer, 2011, vol. 6950.

Books or Proceedings Editing

[37] C. GAVOILLE, P. FRAIGNIAUD (editors). *Proceedings of the 30th Annual ACM Symposium on Principles of Distributed Computing, PODC 2011*, ACM, San Jose, CA, USA, June 2011.

Research Reports

[38] F. BACCELLI, F. MATHIEU, I. NORROS. *Performance of P2P Networks with Spatial Interactions of Peers*, INRIA, August 2011, n^{o} RR-7713, http://hal.inria.fr/inria-00615523/en.

[39] C. GAVOILLE, Q. GODFROY, L. VIENNOT. *Node-Disjoint Multipath Spanners and their Relationship with Fault-Tolerant Spanners*, LaBRI, September 2011, 15 pages, http://hal.archives-ouvertes.fr/hal-00622915/en/.

[40] F. MATHIEU. *On Using Seeders for P2P Live Streaming*, INRIA, 2011, n^{o} RR-7608, http://hal.inria.fr/inria-00588747/en.

Scientific Popularization

[41] C. DELPORTE-GALLET, H. FAUCONNIER. *Objets partagés et détecteurs de défaillances*, in "Technique et Science Informatiques", 2011, vol. 30/7, p. 841-871.