



IN PARTNERSHIP WITH:
CNRS

**Université Denis Diderot
(Paris 7)**

Activity Report 2012

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

1. Members	1
2. Overall Objectives	1
3. Application Domains	2
4. New Results	2
4.1. Understanding graph representations	2
4.1.1. Notions of Connectivity in Overlay Networks	2
4.1.2. Connected graph searching	2
4.1.3. Computing with Large Populations Using Interactions	3
4.1.4. Collaborative Search on the Plane without Communication	3
4.1.5. Memory Lower Bounds for Randomized Collaborative Search and Implications for Biology	4
4.1.6. What Can be Computed without Communications?	4
4.1.7. Decidability Classes for Mobile Agents Computing modularity	4
4.1.8. Randomized Distributed Decision	4
4.1.9. The Worst Case Behavior of Randomized Gossip	5
4.1.10. Asymptotic modularity	5
4.1.11. Modeling social networks	6
4.1.12. Additive Spanners and Distance and Routing Labeling Schemes for Hyperbolic Graphs	6
4.1.13. Constructing a Minimum phylogenetic Network from a Dense triplet Set	6
4.1.14. Algorithms for Some H -Join Decompositions	6
4.1.15. Detecting 2-joins faster	7
4.2. Large Scale Networks Performance and Modeling	7
4.2.1. Spatial Interactions of Peers and Performance of File Sharing Systems	7
4.2.2. User Behavior Modeling: Four Months in DailyMotion	7
4.2.3. Multi-Carrier Networks: on the Manipulability of Voting Systems	7
4.3. Fault Tolerance in Distributed Networks	8
4.3.1. Wait-Freedom with Advice	8
4.3.2. Partial synchrony based on set timeliness	8
4.3.3. Byzantine Agreement with Homonyms in Synchronous Systems	8
4.3.4. Homonyms with forgeable identifiers	9
4.4. Discrete Optimization Algorithms	9
4.4.1. Estimating satisfiability	9
4.4.2. Attractive force search algorithm for piecewise convex maximization problems	9
4.4.3. B-spline interpolation: Toeplitz inverse under corner perturbations	9
5. Bilateral Contracts and Grants with Industry	10
5.1.1. Peer-to-peer for high quality Internet radio	10
5.1.2. Manipulability of voting systems and applications to networks	10
5.1.3. Extension of PageRank for Social Networks	10
6. Partnerships and Cooperations	10
6.1. Regional Initiatives	10
6.2. National Initiatives	10
6.2.1. ANR SONGS	10
6.2.2. ANR Prose	10
6.2.3. ANR Shaman	11
6.2.4. ANR Displexity	11
6.3. European Initiatives	12
6.4. International Initiatives	13
7. Dissemination	13
7.1.1. Teaching	13

7.1.2. Supervision	14
7.1.3. Juries	14
8. Bibliography	14

Project-Team GANG

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

Creation of the Project-Team: July 01, 2007 .

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]
Fabien Mathieu [Senior Researcher Inria, HdR]

Faculty Members

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

PhD Students

François Durand [PhD Student]
The-Dang Huynh [PhD Student]

Administrative Assistant

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

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. *Notions of Connectivity in Overlay Networks*

Participants: Yuval Emek, Pierre Fraigniaud, Amos Korman, Shay Kutten, David Peleg.

How well connected is the network? This is one of the most fundamental questions one would ask when facing the challenge of designing a communication network. Three major notions of connectivity have been considered in the literature, but in the context of traditional (single-layer) networks, they turn out to be equivalent. The paper [17], introduces a model for studying the three notions of connectivity in multi-layer networks. Using this model, it is easy to demonstrate that in multi-layer networks the three notions may differ dramatically. Unfortunately, in contrast to the single-layer case, where the values of the three connectivity notions can be computed efficiently, it has been recently shown in the context of WDM networks (results that can be easily translated to our model) that the values of two of these notions of connectivity are hard to compute or even approximate in multi-layer networks. The current paper shed some positive light into the multi-layer connectivity topic: we show that the value of the third connectivity notion can be computed in polynomial time and develop an approximation for the construction of well connected overlay networks.

4.1.2. *Connected graph searching*

Participants: Lali Barrière, Paola Flocchini, Fedor V. Fomin, Pierre Fraigniaud, Nicolas Nisse, Nicola Santoro, Dimitrios M. Thilikos.

In the graph searching game the opponents are a set of searchers and a fugitive in a graph. The searchers try to capture the fugitive by applying some sequence of moves that include placement, removal, or sliding of a searcher along an edge. The fugitive tries to avoid capture by moving along unguarded paths. The search number of a graph is the minimum number of searchers required to guarantee the capture of the fugitive. In [2], we initiate the study of this game under the natural restriction of connectivity where we demand that in each step of the search the locations of the graph that are clean (i.e. non-accessible to the fugitive) remain connected. We give evidence that many of the standard mathematical tools used so far in classic graph searching fail under the connectivity requirement. We also settle the question on “the price of connectivity”, that is, how many searchers more are required for searching a graph when the connectivity demand is imposed. We make estimations of the price of connectivity on general graphs and we provide tight bounds for the case of trees. In particular, for an n -vertex graph the ratio between the connected searching number and the non-connected one is while for trees this ratio is always at most 2. We also conjecture that this constant-ratio upper bound for trees holds also for all graphs. Our combinatorial results imply a complete characterization of connected graph searching on trees. It is based on a forbidden-graph characterization of the connected search number. We prove that the connected search game is monotone for trees, i.e. restricting search strategies to only those where the clean territories increase monotonically does not require more searchers. A consequence of our results is that the connected search number can be computed in polynomial time on trees, moreover, we show how to make this algorithm distributed. Finally, we reveal connections of this parameter to other invariants on trees such as the Horton–Strahler number.

4.1.3. Computing with Large Populations Using Interactions

Participants: Olivier Bournez, Pierre Fraigniaud, Xavier Koegler.

We define in [12], a general model capturing the behavior of a population of anonymous agents that interact in pairs. This model captures some of the main features of opportunistic networks, in which nodes (such as the ones of a mobile ad hoc networks) meet sporadically. For its reminiscence to Population Protocol, we call our model Large-Population Protocol, or LPP. We are interested in the design of LPPs enforcing, for every $\nu \in [0, 1]$, a proportion ν of the agents to be in a specific subset of marked states, when the size of the population grows to infinity; In which case, we say that the protocol computes ν . We prove that, for every $\nu \in [0, 1]$, ν is computable by a LPP if and only if ν is algebraic. Our positive result is constructive. That is, we show how to construct, for every algebraic number $\nu \in [0, 1]$, a protocol which computes ν .

4.1.4. Collaborative Search on the Plane without Communication

Participants: Ofer Feinerman, Zvi Lotker, Amos Korman, Jean-Sébastien Sereni.

In [19], we use distributed computing tools to provide a new perspective on the behavior of cooperative biological ensembles. We introduce the Ants Nearby Treasure Search (ANTS) problem, a generalization of the classical cow-path problem which is relevant for collective foraging in animal groups. In the ANTS problem, k identical (probabilistic) agents, initially placed at some central location, collectively search for a treasure in the two-dimensional plane. The treasure is placed at a target location by an adversary and the goal is to find it as fast as possible as a function of both k and D , where D is the distance between the central location and the target. This is biologically motivated by cooperative, central place foraging, such as performed by ants around their nest. In this type of search there is a strong preference to locate nearby food sources before those that are further away. We focus on trying to find what can be achieved if communication is limited or altogether absent. Indeed, to avoid overlaps agents must be highly dispersed making communication difficult. Furthermore, if the agents do not commence the search in synchrony, then even initial communication is problematic. This holds, in particular, with respect to the question of whether the agents can communicate and conclude their total number, k . It turns out that the knowledge of k by the individual agents is crucial for performance. Indeed, it is a straightforward observation that the time required for finding the treasure is $\Omega(D + D^2/k)$, and we show in this paper that this bound can be matched if the agents have knowledge of k up to some constant approximation. We present a tight bound for the competitive penalty that must be paid, in the running time, if the agents have no information about k . Specifically, this bound is slightly more than logarithmic in the number of agents. In addition, we give a lower bound for the setting in which the agents are given some estimation of

k . Informally, our results imply that the agents can potentially perform well without any knowledge of their total number k , however, to further improve, they must use some information regarding k . Finally, we propose a uniform algorithm that is both efficient and extremely simple, suggesting its relevance for actual biological scenarios.

4.1.5. *Memory Lower Bounds for Randomized Collaborative Search and Implications for Biology*

Participants: Ofer Feinerman, Amos Korman.

Initial knowledge regarding group size can be crucial for collective performance. We study in [18], this relation in the context of the Ants Nearby Treasure Search (ANTS) problem, which models natural cooperative foraging behavior such as that performed by ants around their nest. In this problem, k (probabilistic) agents, initially placed at some central location, collectively search for a treasure on the two-dimensional grid. The treasure is placed at a target location by an adversary and the goal is to find it as fast as possible as a function of both k and D , where D is the (unknown) distance between the central location and the target. It is easy to see that $T = \Omega(D + D^2/k)$ time units are necessary for finding the treasure. Recently, it has been established that $O(T)$ time is sufficient if the agents know their total number k (or a constant approximation of it), and enough memory bits are available at their disposal. In this paper, we establish lower bounds on the agent memory size required for achieving certain running time performances. To the best of our knowledge, these bounds are the first non-trivial lower bounds for the memory size of probabilistic searchers. For example, for every given positive constant ϵ , terminating the search by time $O(\log^{1-\epsilon} k \cdot T)$ requires agents to use $\Omega(\log \log k)$ memory bits.

From a high level perspective, we illustrate how methods from distributed computing can be useful in generating lower bounds for cooperative biological ensembles. Indeed, if experiments that comply with our setting reveal that the ants' search is time efficient, then our theoretical lower bounds can provide some insight on the memory they use for this task.

4.1.6. *What Can be Computed without Communications?*

Participants: Heger Arfaoui, Pierre Fraigniaud.

When playing the boolean game (δ, f) , two players, upon reception of respective inputs x and y , must respectively output a and b satisfying $\delta(a, b) = f(x, y)$, in absence of any communication. It is known that, for $\delta(a, b) = a \oplus b$, the ability for the players to use entangled quantum bits (qbits) helps. In [10], we show that, for δ different from the exclusive-or operator, quantum correlations do not help. This result is an invitation to revisit the theory of distributed checking, a.k.a. distributed verification, currently stucked to the usage of decision functions δ based on the and-operator, hence potentially preventing us from using the potential benefit of quantum effects.

4.1.7. *Decidability Classes for Mobile Agents Computing modularity*

Participants: Andrzej Pelc, Pierre Fraigniaud.

We establish in [21], a classification of decision problems that are to be solved by mobile agents operating in unlabeled graphs, using a deterministic protocol. The classification is with respect to the ability of a team of agents to solve the problem, possibly with the aid of additional information. In particular, our focus is on studying differences between the decidability of a decision problem by agents and its verifiability when a certificate for a positive answer is provided to the agents. Our main result shows that there exists a natural complete problem for mobile agent verification. We also show that, for a single agent, three natural oracles yield a strictly increasing chain of relative decidability classes.

4.1.8. *Randomized Distributed Decision*

Participants: Pierre Fraigniaud, Amos Korman, Merav Parter, David Peleg.

The paper [20] tackles the power of randomization in the context of locality by analyzing the ability to “boost” the success probability of deciding a distributed language. The main outcome of this analysis is that the distributed computing setting contrasts significantly with the sequential one as far as randomization is concerned. Indeed, we prove that in some cases, the ability to increase the success probability for deciding distributed languages is rather limited.

We focus on the notion of a (p, q) -decider for a language L , which is a distributed randomized algorithm that accepts instances in L with probability at least p and rejects instances outside of L with probability at least q . It is known that every hereditary language that can be decided in t rounds by a (p, q) -decider, where $p^2 + q > 1$, can be decided deterministically in $O(t)$ rounds. One of our results gives evidence supporting the conjecture that the above statement holds for all distributed languages and not only for hereditary ones, by proving the conjecture for the restricted case of path topologies. For the range below the aforementioned threshold, namely, $p^2 + q \leq 1$, we study the class $B_k(t)$ (for $k \in \mathbb{N}^* \cup \{\infty\}$) of all languages decidable in at most t rounds by a (p, q) -decider, where $p^{1+\frac{1}{k}} + q > 1$. Since every language is decidable (in zero rounds) by a (p, q) -decider satisfying $p + q = 1$, the hierarchy B_k provides a spectrum of complexity classes between determinism ($k = 1$, under the above conjecture) and complete randomization ($k = \infty$). We prove that all these classes are separated, in a strong sense: for every integer $k \geq 1$, there exists a language L satisfying $L \in B_{k+1}(0)$ but $L \notin B_k(t)$ for any $t = o(n)$. In addition, we show that $B_\infty(t)$ does not contain all languages, for any $t = o(n)$. In other words, we obtain the hierarchy $B(t) \subset B_2(t) \subset \dots \subset B_\infty(t) \subset \text{All}$. Finally, we show that if the inputs can be restricted in certain ways, then the ability to boost the success probability becomes almost null, and in particular, derandomization is not possible even beyond the threshold $p^2 + q = 1$.

4.1.9. The Worst Case Behavior of Randomized Gossip

Participants: Hervé Baumann, Pierre Fraigniaud, Hovhannes A. Harutyunyan, Rémi de Verclos.

In [11] we consider the quasi-random rumor spreading model introduced by Doerr, Friedrich, and Sauerwald in [SODA 2008], hereafter referred to as the list-based model. Each node is provided with a cyclic list of all its neighbors, chooses a random position in its list, and from then on calls its neighbors in the order of the list. This model is known to perform asymptotically at least as well as the random phone-call model, for many network classes. Motivated by potential applications of the list-based model to live streaming, we are interested in its worst case behavior.

Our first main result is the design of an $O(m + n \log n)$ -time algorithm that, given any n -node m -edge network G , and any source-target pair $s, t \in V(G)$, computes the maximum number of rounds it may take for a rumor to be broadcast from s to t in G , in the list-based model. This algorithm yields an $O(n(m + n \log n))$ -time algorithm that, given any network G , computes the maximum number of rounds it may take for a rumor to be broadcast from any source to any target, in the list-based model. Hence, the list-based model is computationally easy to tackle in its basic version.

The situation is radically different when one is considering variants of the model in which nodes are aware of the status of their neighbors, i.e., are aware of whether or not they have already received the rumor, at any point in time. Indeed, our second main result states that, unless $P=NP$, the worst case behavior of the list-based model with the additional feature that every node is perpetually aware of which of its neighbors have already received the rumor cannot be approximated in polynomial time within a $(\frac{1}{n})^{\frac{1}{2}-\epsilon}$ multiplicative factor, for any $\epsilon > 0$. As a byproduct of this latter result, we can show that, unless $P=NP$, there are no PTAS enabling to approximate the worst case behavior of the list-based model, whenever every node perpetually keeps track of the subset of its neighbors which have sent the rumor to it so far.

4.1.10. Asymptotic modularity

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

Modularity (Newman-Girvan) has been introduced as a quality measure for graph partitioning. 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 the modularity of a variety of graph classes. In [23], we first consider simple graph classes such as tori and hypercubes. We show that these regular graph families

have asymptotic modularity 1 (that is the maximum possible). We extend this result to trees with bounded degree, allowing us to give a lower bound of 2 over average degree for graph classes with low maximum degree (included power law graphs for a sufficiently large exponent).

4.1.11. Modeling social networks

Participants: Nidhi Hegde, Laurent Massoulié, Laurent Viennot.

Social networks offer users new means of accessing information, essentially relying on “social filtering”, i.e. propagation and filtering of information by social contacts. The sheer amount of data flowing in these networks, combined with the limited budget of attention of each user, makes it difficult to ensure that social filtering brings relevant content to the interested users. Our motivation in [26] is to measure to what extent self-organization of the social network results in efficient social filtering. To this end we introduce flow games, a simple abstraction that models network formation under selfish user dynamics, featuring user-specific interests and budget of attention. In the context of homogeneous user interests, we show that selfish dynamics converge to a stable network structure (namely a pure Nash equilibrium) with close-to-optimal information dissemination. We show in contrast, for the more realistic case of heterogeneous interests, that convergence, if it occurs, may lead to information dissemination that can be arbitrarily inefficient, as captured by an unbounded “price of anarchy”. Nevertheless the situation differs when users’ interests exhibit a particular structure, captured by a metric space with low doubling dimension. In that case, natural autonomous dynamics converge to a stable configuration. Moreover, users obtain all the information of interest to them in the corresponding dissemination, provided their budget of attention is logarithmic in the size of their interest set.

4.1.12. Additive Spanners and Distance and Routing Labeling Schemes for Hyperbolic Graphs

Participants: Victor Chepoi, Feodor Dragan, Bertrand Estellon, Michel Habib, Yann Vaxès, Yang Xiang.

δ -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], we study unweighted δ -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 construction 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 δ -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.

4.1.13. Constructing a Minimum phylogenetic Network from a Dense triplet Set

Participants: Michel Habib, Thu-Hien To.

For a given set \mathcal{L} of species and a set \mathcal{T} of triplets on \mathcal{L} , we seek to construct a phylogenetic network which is consistent with \mathcal{T} i.e. which represents all triplets of \mathcal{T} . The level of a network is defined as the maximum number of hybrid vertices in its biconnected components. When \mathcal{T} is dense, there exist polynomial time algorithms to construct level-0, 1 and 2 networks (Aho et al., 1981; Jansson, Nguyen and Sung, 2006; Jansson and Sung, 2006; Iersel et al., 2009). For higher levels, partial answers were obtained in the paper by Iersel and Kelk (2008), with a polynomial time algorithm for simple networks. In [9] this paper, we detail the first complete answer for the general case, solving a problem proposed in Jansson and Sung (2006) and Iersel et al. (2009). For any k fixed, it is possible to construct a level- k network having the minimum number of hybrid vertices and consistent with \mathcal{T} , if there is any, in time $O\left(|\mathcal{T}|^{k+1} n^{\lfloor \frac{4k}{3} \rfloor}\right)$.

4.1.14. Algorithms for Some H -Join Decompositions

Participants: Michel Habib, Antoine Mamcarz, Fabien de Montgolfier.

A homogeneous pair (also known as a 2-module) of a graph is a pair $\{M_1, M_2\}$ of disjoint vertex subsets such that for every vertex $x \notin (M_1 \cup M_2)$ and $i \in \{1, 2\}$, x is either adjacent to all vertices in M_i or to none of them. First used in the context of perfect graphs [Chvátal and Sbihi 1987], it is a generalization of splits (a.k.a 1-joins) and of modules. The algorithmics to compute them appears quite involved. In [22], we describe an $O(mn^2)$ -time algorithm computing (if any) a homogeneous pair, which not only improves a previous bound of $O(mn^3)$ [Everett, Klein and Reed 1997], but also uses a nice structural property of homogenous pairs. Our result can be extended to compute the whole homogeneous pair decomposition tree, within the same complexity. Using similar ideas, we present an $O(nm^2)$ -time algorithm to compute a N -join decomposition of a graph, improving a previous $O(n^6)$ algorithm [Feder et al. 2005]. These two decompositions are special case of H -joins [Bui-Xuan, Telle and Vatshelle 2010] to which our techniques apply.

4.1.15. Detecting 2-joins faster

Participants: Pierre Charbit, Michel Habib, Nicolas Trotignon, Kristina Vušković.

2-joins are edge cutsets that naturally appear in the decomposition of several classes of graphs closed under taking induced subgraphs, such as balanced bipartite graphs, even-hole-free graphs, perfect graphs and claw-free graphs. Their detection is needed in several algorithms, and is the slowest step for some of them. The classical method to detect a 2-join takes $O(n^3m)$ time where n is the number of vertices of the input graph and m the number of its edges. To detect *non-path* 2-joins (special kinds of 2-joins that are needed in all of the known algorithms that use 2-joins), the fastest known method takes time $O(n^4m)$. Here, we give an $O(n^2m)$ -time algorithm for both of these problems. A consequence is a speed up of several known algorithms.

4.2. Large Scale Networks Performance and Modeling

4.2.1. Spatial Interactions of Peers and Performance of File Sharing Systems

Participants: François Baccelli, Fabien Mathieu, Ilkka Norros.

We propose in [24] a new model for peer-to-peer networking which takes the network bottlenecks into account beyond the access. This model allows one to cope with key features of P2P networking like degree or locality constraints or the fact that distant peers often have a smaller rate than nearby peers. We show that the spatial point process describing peers in their steady state then exhibits an interesting repulsion phenomenon. We analyze two asymptotic regimes of the peer-to-peer network: the fluid regime and the hard-core regime. We get closed form expressions for the mean (and in some cases the law) of the peer latency and the download rate obtained by a peer as well as for the spatial density of peers in the steady state of each regime, as well as an accurate approximation that holds for all regimes. The analytical results are based on a mix of mathematical analysis and dimensional analysis and have important design implications. The first of them is the existence of a setting where the equilibrium mean latency is a decreasing function of the load, a phenomenon that we call super-scalability.

4.2.2. User Behavior Modeling: Four Months in DailyMotion

Participants: Yannick Carlinet, The Dang Huynh, Bruno Kauffmann, Fabien Mathieu, Ludovic Noirie, Sébastien Tixeuil.

The growth of User-Generated Content (UGC) traffic makes the understanding of its nature a priority for network operators, content providers and equipment suppliers. In [13], we study a four-month dataset that logs all video requests to DailyMotion made by a fixed subset of users. We were able to infer user sessions from raw data, to propose a Markovian model of these sessions, and to study video popularity and its evolution over time. The presented results are a first step for synthesizing an artificial (but realistic) traffic that could be used in simulations or experimental testbeds.

4.2.3. Multi-Carrier Networks: on the Manipulability of Voting Systems

Participants: François Durand, Fabien Mathieu, Ludovic Noirie.

Today, Internet involves many actors who are making revenues on it (operators, companies, service providers,...). It is therefore important to be able to make fair decisions in this large-scale and highly competitive economical ecosystem. One of the main issues is to prevent actors from manipulating the natural outcome of the decision process. For that purpose, game theory is a natural framework. In that context, voting systems represent an interesting alternative that, to our knowledge, has not yet been considered. They allow competing entities to decide among different options. Strong theoretical results showed that all voting systems are susceptible to be manipulated by one single voter, except for some "degenerated" and non-acceptable cases. However, very little is known about how much a voting system is manipulable in practical scenarios. In [25], we investigate empirically the use of voting systems for choosing end-to-end paths in multi-carrier networks, analyzing their manipulability and their economical efficiency. We show that one particular system, called Single Transferable Vote (STV), is largely more resistant to manipulability than the natural system which tries to get the economical optimum. Moreover, STV manages to select paths close to the economical optimum, whether the participants try to cheat or not.

4.3. Fault Tolerance in Distributed Networks

4.3.1. Wait-Freedom with Advice

Participants: Carole Delporte-Gallet, Hugues Fauconnier, Eli Gafni, Petr Kuznetsov.

In [14], we motivate and propose a new way of thinking about failure detectors which allows us to define, quite surprisingly, what it means to solve a distributed task *wait-free using a failure detector*. In our model, the system is composed of *computation* processes that obtain inputs and are supposed to output in a finite number of steps and *synchronization* processes that are subject to failures and can query a failure detector. We assume that, under the condition that *correct* synchronization processes take sufficiently many steps, they provide the computation processes with enough *advice* to solve the given task wait-free: every computation process outputs in a finite number of its own steps, regardless of the behavior of other computation processes. Every task can thus be characterized by the *weakest* failure detector that allows for solving it, and we show that every such failure detector captures a form of set agreement. We then obtain a complete classification of tasks, including ones that evaded comprehensible characterization so far, such as renaming or weak symmetry breaking.

4.3.2. Partial synchrony based on set timeliness

Participants: Markos Aguilera, Carole Delporte-Gallet, Hugues Fauconnier, Sam Toueg.

We introduce in [1], a new model of partial synchrony for read-write shared memory systems. This model is based on the simple notion of set timeliness—a natural generalization of the seminal concept of timeliness in the partially synchrony model of Dwork et al. (J. ACM 35(2):288–323, 1988). Despite its simplicity, the concept of set timeliness is powerful enough to define a family of partially synchronous systems that closely match individual instances of the t -resilient k -set agreement problem among n processes, henceforth denoted (t, k, n) -agreement. In particular, we use it to give a partially synchronous system that is synchronous enough for solving (t, k, n) -agreement, but not enough for solving two incrementally stronger problems, namely, $(t + 1, k, n)$ -agreement, which has a slightly stronger resiliency requirement, and $(t, k - 1, n)$ -agreement, which has a slightly stronger agreement requirement. This is the first partially synchronous system that separates these sub-consensus problems. The above results show that set timeliness can be used to study and compare the partial synchrony requirements of problems that are strictly weaker than consensus.

4.3.3. Byzantine Agreement with Homonyms in Synchronous Systems

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

In [15], [6], we consider the Byzantine agreement problem (BA) in synchronous systems with homonyms. In this model different processes may have the same authenticated identifier. In such a system of n processes sharing a set of l identifiers, we define a distribution of the identifiers as an integer partition of n into l parts n_1, \dots, n_l giving for each identifier i the number of processes having this identifier.

Assuming that the processes know the distribution of identifiers we give a necessary and sufficient condition on the integer partition of n to solve the Byzantine agreement with at most t Byzantine processes. Moreover we prove that there exists a distribution of l identifiers enabling to solve Byzantine agreement with at most t Byzantine processes if and only if $n > 3t$, $l > t$ and where $r = n \bmod l$.

This bound is to be compared with the $l > 3t$ bound proved in Delporte-Gallet et al. (2011) when the processes do not know the distribution of identifiers.

4.3.4. *Homonyms with forgeable identifiers*

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

In [16], we refine the Byzantine Agreement problem (BA) in synchronous systems with homonyms, in the particular case where some identifiers may be forgeable. More precisely, the n processes share a set of l ($1 \leq l \leq n$) identifiers. Assuming that at most t processes may be Byzantine and at most k ($t \leq k \leq l$) of these identifiers are forgeable in the sense that any Byzantine process can falsely use them, we prove that Byzantine Agreement problem is solvable if and only if $l > 2t + k$. Moreover we extend this result to systems with authentication by signatures in which at most k signatures are forgeable and we prove that Byzantine Agreement problem is solvable if and only if $l > t + k$.

4.4. Discrete Optimization Algorithms

4.4.1. *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 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 [3], 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.

4.4.2. *Attractive force search algorithm for piecewise convex maximization problems*

Participants: Dominique Fortin, Ider Tseveendorj.

In [8], 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.

4.4.3. *B-spline interpolation: Toeplitz inverse under corner perturbations*

Participant: Dominique Fortin.

For Toeplitz matrices associated with degree 3 and 4 uniform B-spline interpolation, the inverse may be analytically known [7], saving the standard inverse calculations. It generalizes to any degree as a row of the Eulerian numbers triangle.

5. Bilateral Contracts and Grants with Industry

5.1. Bilateral Contracts with Industry

5.1.1. *Peer-to-peer for high quality Internet radio*

Participant: Fabien Mathieu.

A contract has been signed between Inria, RadioCeros and the ARITT Center. Gang is to provide a feasibility study on the subject of the use of Peer-to-peer mechanisms for high quality Internet radio.

5.1.2. *Manipulability of voting systems and applications to networks*

Participants: Fabien Mathieu, François Durand.

Alcatel grants ADR LINCS to study applicability of voting systems to loosely connected networks (Peer-to-peer, social networks...).

5.1.3. *Extension of PageRank for Social Networks*

Participants: Fabien Mathieu, The-Dang Huynh.

ALCATEL is funding a CIFRE PhD for carrying PageRank techniques to Social Networks.

6. Partnerships and Cooperations

6.1. Regional Initiatives

6.1.1. *PEFICAMO*

Participants: Carole Delporte-Gallet, Hugues Fauconnier, 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 SONGS*

Participant: Fabien Mathieu.

The goal of the SONGS project is to extend the applicability of the SimGrid simulation framework from Grids and Peer-to-peer systems to Clouds and High Performance Computation systems. Each type of large-scale computing system will be addressed through a set of use cases and lead by researchers recognized as experts in this area.

6.2.2. *ANR Prose*

Participants: Pierre Fraigniaud, Amos Korman, Laurent Viennot.

Managed by University Paris Diderot, P. Fraigniaud.

Online social networks are among the most popular sites on the Web and continue to grow rapidly. They provide mechanisms to establish identities, share content and information, and create relationships. With the emergence of a new generation of powerful mobile devices that enable wireless ad hoc communication, it is time to extend social networking to the mobile world. Such an ad hoc social networking environment is full of opportunities. As opposed to the use of personal computers, a mobile phone is a strictly personal device, always on, with several wireless interfaces that include a short range communication with nearby nodes. Applications such as notification of status updates, sharing of user generated content, documents tagging, rating/recommendation and bookkeeping can be deployed “on the move” on top of contacts established through short range communication. It requires to deploy social networking applications in a delay tolerant manner using opportunistic social contacts as in a peer to peer network, as well as new advanced content recommendation engines.

The Prose project is a collective and multi-disciplinary effort to design opportunistic contact sharing schemes, and characterizes the environmental conditions, the usage constraint, as well as the algorithmic and architecture principles that let them operate. The partners of the Prose project will engage in this exploration through various expertise: network measurement, traffic monitoring from a real application, system design, behavioral study, analysis of distributed algorithms, theory of dynamic graph, networking modeling, and performance evaluation. As part of this project, the partners will be involved in the analysis of the content received and accessed by users of a real commercial application (PlayAdz), and will participate to the design of a new promotion advertisement service.

6.2.3. ANR Shaman

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

SHAMAN (Self-organizing and Healing Architectures for Malicious and Adversarial Networks) is an ANR VERSO Project (2009-2012).

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

SHAMAN focuses on the algorithmic foundations of resource-constrained autonomous large scale systems, dedicated to enabling the sustainability of network functions in spite of abrupt system evolutions, component failures, and attacks. We foresee original solutions in the general frameworks of self-stabilization, failure detection, and robust protocols. Our first objective is the design of obligate but realistic models encompassing anonymity, dynamism, and/or malicious behavior. Our second objective is to evaluate both the theoretical power, and the practical functionality, of these models, by confronting them to their ability of designing efficient algorithms and protocols for dynamic and malicious environments. This evaluation will be tackled in two complementary application domains: wireless sensor networks, and peer-to-peer systems. The primary outcome of SHAMAN should be the demonstration of reliable middleware bricks that could be integrated in real distributed platforms.

6.2.4. ANR Displexity

Participants: Carole Delporte-Gallet, Hugues Fauconnier, Pierre Fraigniaud, Arfoui Heger, 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.

Distributed computation keep raising new questions concerning computability and complexity. For instance, as far as fault-tolerant distributed computing is concerned, impossibility results do not depend on the computational power of the processes, demonstrating a form of undecidability which is significantly different from the one encountered in sequential computing. In the same way, as far as network computing is concerned, the impossibility of solving certain tasks locally does not depend on the computational power of the individual processes.

The main goal of DISPLEXITY (for DIStributed computing: computability and COMPLEXITY) is to establish the scientific foundations for building up a consistent theory of computability and complexity for distributed computing.

One difficulty to be faced by DISPLEXITY is to reconcile the different sub-communities corresponding to a variety of classes of distributed computing models. The current distributed computing community may indeed be viewed as two not necessarily disjoint sub-communities, one focusing on the impact of temporal issues, while the other focusing on the impact of spatial issues. The different working frameworks tackled by these two communities induce different objectives: computability is the main concern of the former, while complexity is the main concern of the latter.

Within DISPLEXITY, the reconciliation between the two communities will be achieved by focusing on the same class of problems, those for which the distributed outputs are interpreted as a single binary output: yes or no. Those are known as the yes/no-problems. The strength of DISPLEXITY is to gather specialists of the two main streams of distributed computing. Hence, DISPLEXITY will take advantage of the experience gained over the last decade by both communities concerning the challenges to be faced when building up a complexity theory encompassing more than a fragment of the field.

In order to reach its objectives, DISPLEXITY aims at achieving the following tasks:

- Formalizing yes/no-problems (decision problems) in the context of distributed computing. Such problems are expected to play an analogous role in the field of distributed computing as that played by decision problems in the context of sequential computing.
- Formalizing decision problems (yes/no-problems) in the context of distributed computing. Such problems are expected to play an analogous role in the field of distributed computing as that played by decision problems in the context of sequential computing.
- Revisiting the various explicit (e.g., failure-detectors) or implicit (e.g., a priori information) notions of oracles used in the context of distributed computing allowing us to express them in terms of decidability/complexity classes based on oracles.
- Identifying the impact of non-determinism on complexity in distributed computing. In particular, DISPLEXITY aims at a better understanding of the apparent lack of impact of non-determinism in the context of fault-tolerant computing, to be contrasted with the apparent huge impact of non-determinism in the context of network computing. Also, it is foreseen that non-determinism will enable the comparison of complexity classes defined in the context of fault-tolerance with complexity classes defined in the context of network computing.
- Last but not least, DISPLEXITY will focus on new computational paradigms and frameworks, including, but not limited to distributed quantum computing and algorithmic game theory (e.g., network formation games).

The project will have to face and solve a number of challenging problems. Hence, we have built the DISPLEXITY consortium so as to coordinate the efforts of those worldwide leaders in Distributed Computing who are working in our country. A successful execution of the project will result in a tremendous increase in the current knowledge and understanding of decentralized computing and place us in a unique position in the field.

6.3. European Initiatives

6.3.1. FP7 Projects

6.3.1.1. 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 Politècnica 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.

Description: In this collaboration, we mainly 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.

6.4. International Initiatives

6.4.1. Internet Technologies and Architectures

Participant: Fabien Mathieu.

The aim of this project is to build a community of researchers focusing on fundamental theoretical issues of future networking, including such topics as communication theory, network information theory, distributed algorithms, self-organization and game theory, modeling of large random and complex networks and structures. Partners Inria, VTT, Aalto University, Eindhoven University are gathered under EIT ICT Labs Project Fundamentals of Networking (FUN). <http://eit.ictlabs.eu/ict-labs/all-events/article/fundamentals-of-future-networking-workshop/>.

7. Dissemination

7.1. Teaching - Supervision - Juries

7.1.1. Teaching

Master MPRI University of Paris Diderot:

- M. Habib, graph algorithms, 12 hours;
- P. Fraigniaud, "Algorithmique distribuée pour les réseaux", 12 hours;
- C. Delporte and H. Fauconnier, "Algorithmique distribuée avec mémoire partagée"; C. Delporte, 12 hours and H. Fauconnier, 12 hours;
- L. Viennot, "Structures de données distribuées et routage", 12 hours.

Master Professional University of Paris Diderot:

- M. Habib, Search Engines, 50 hours;
- M. Habib, Parallelism and mobility which includes peer-to-peer overlay networks, 50 hours;
- C. Delporte, Distributed programming, 85 hours;
- H. Fauconnier, Internet Protocols and Distributed algorithms, 85 hours.

Master: F. Mathieu, Peer-to-Peer Techniques, 30 hours, University of Paris 6;

D.U.T. : Y. Boufkhad, computer science and networks, 192 hours, University of Paris Diderot;

U.F.R.: F. de Montgolfier, foundation of computer science, algorithmics, and computer architecture, 192 hours, University of Paris Diderot;

Master: F. de Montgolfier, Peer-to-Peer theory and application, M2, University of Marne-la-Vallée.

7.1.2. Supervision

PhD : Hervé Baumann, "Diffusion décentralisée d'information dans les systèmes distribués", University of Paris Diderot, september 24, 2012

Xavier Koegler, "Protocoles de population, jeux et grandes populations", University of Paris Diderot, september 13, 2012

PhD in progress : François Durand, "Manipulabilité des systèmes de vote et applications aux réseaux", since 2012, supervised by Fabien Mathieu and Ludovic Noirie

Jérémie Dusart, "Parcours de graphes de cocomparabilité", since 2011, supervised by Michel Habib
Antoine Mamcarz, "Algorithmes de décomposition de graphes", since 2010, supervised by Michel Habib

The-Dang Huynh , "Extensions de PageRank et Applications aux Réseaux Sociaux", since 2012, supervised by Fabien Mathieu and Dohy Hong

Hung Tran, Failure detection with Byzantine adversary, since 2010, supervised by Hugues Fauconnier and Carole Delporte

7.1.3. Juries

HdR review: L. Viennot reviews Arnaud Legout's HDR thesis on "Efficacité et vie privée : de BitTorrent à Skype";

PhD review: L. Viennot reviews Bio Mikaila Toko Worou's PhD thesis on "Outils algorithmiques de détection des communautés dans les réseaux";

PhD thesis : L. Viennot was examiner of Quentin Godfroy's thesis on "From spanners to multipath spanners"

F. Mathieu was examiner of Claudio Testa's thesis "On the congestion control of Peer-to-peer applications: the LEDBAT case".

8. Bibliography

Publications of the year

Articles in International Peer-Reviewed Journals

- [1] M. K. AGUILERA, C. DELPORTE-GALLET, H. FAUCCONNIER, S. TOUEG. *Partial synchrony based on set timeliness*, in "Distributed Computing", 2012, vol. 25, n^o 3, p. 249-260.

- [2] L. BARRIÈRE, P. FLOCCHINI, F. V. FOMIN, P. FRAIGNIAUD, N. NISSE, N. SANTORO, D. M. THILIKOS. *Connected graph searching*, in "Inf. Comput.", 2012, vol. 219, p. 1-16.
- [3] Y. BOUFGHAD, T. HUGEL. *Estimating satisfiability*, in "Discrete Applied Mathematics", 2012, vol. 160, n^o 1-2, p. 61-80.
- [4] P. CHARBIT, M. HABIB, N. TROTIGNON, K. VUŠKOVIĆ. *Detecting 2-joins faster*, in "J. Discrete Algorithms", 2012, vol. 17, p. 60–66, <http://dx.doi.org/10.1016/j.jda.2012.11.003>.
- [5] V. CHEPOI, F. F. DRAGAN, B. ESTELLON, M. HABIB, Y. VAXÈS, Y. XIANG. *Additive Spanners and Distance and Routing Labeling Schemes for Hyperbolic Graphs*, in "Algorithmica", 2012, vol. 62, n^o 3-4, p. 713-732.
- [6] C. DELPORTE-GALLET, H. FAUCONNIER, H. TRAN-THE. *Byzantine agreement with homonyms in synchronous systems*, in "Theoretical Computer Science", 2012, n^o 9129.
- [7] D. FORTIN. *B-spline Toeplitz inverse under corner perturbations*, in "IJPAM", 2012, vol. 77, n^o 1, p. 107–118.
- [8] D. FORTIN, I. TSEVEENDORJ. *Attractive force search algorithm for piecewise convex maximization problems*, in "Optim. Lett.", 2012, vol. 6, n^o 7, p. 1317–1333, <http://dx.doi.org/10.1007/s11590-011-0395-y>.
- [9] M. HABIB, T.-H. TO. *Constructing a Minimum phylogenetic Network from a Dense triplet Set*, in "J. Bioinformatics and Computational Biology", 2012, vol. 10, n^o 5.

International Conferences with Proceedings

- [10] H. ARFAOUI, P. FRAIGNIAUD. *Brief announcement: what can be computed without communication?*, in "PODC", 2012, p. 87-88.
- [11] H. BAUMANN, P. FRAIGNIAUD, H. A. HARUTYUNYAN, R. DE VERCLOS. *The Worst Case Behavior of Randomized Gossip*, in "TAMC", 2012, p. 330-345.
- [12] O. BOURNEZ, P. FRAIGNIAUD, X. KOEGLER. *Computing with Large Populations Using Interactions*, in "MFCS", 2012, p. 234-246.
- [13] Y. CARLINET, T. D. HUYNH, B. KAUFFMANN, F. MATHIEU, L. NOIRIE, S. TIXEUIL. *Four Months in DailyMotion: Dissecting User Video Requests*, in "TRAC 2012 - 3rd International Workshop on TRaffic Analysis and Classification", Limassol, Cyprus, August 2012, <http://hal.inria.fr/hal-00692095>.
- [14] C. DELPORTE-GALLET, H. FAUCONNIER, E. GAFNI, P. KUZNETSOV. *Wait-Freedom with Advice.*, in "Proceedings of the Thirty First Annual ACM Symposium on Principles of Distributed Computing, PODC 2012", Madeira, Portugal, ACM, July 16-18 2012.
- [15] C. DELPORTE-GALLET, H. FAUCONNIER, H. TRAN-THE. *Byzantine Agreement with Homonyms in Synchronous Systems*, in "Distributed Computing and Networking - 13th International Conference, ICDCN 2012", Hong Kong, China, L. BONONI, A. K. DATTA, S. DEVISMES, A. MISRA (editors), Lecture Notes in Computer Science, Springer, January 3-6 2012, vol. 7129, p. 76-90.

- [16] C. DELPORTE-GALLET, H. FAUCONNIER, H. TRAN-THE. *Homonyms with forgeable identifiers*, in "19th International Colloquium on Structural Information and Communication Complexity (SIROCCO)", Springer, LNCS, 2012.
- [17] Y. EMEK, P. FRAIGNIAUD, A. KORMAN, S. KUTTEN, D. PELEG. *Notions of Connectivity in Overlay Networks*, in "Proceedings of the 19th International Colloquium on Structural Information and Communication Complexity (SIROCCO 2012)", 2012, p. 25–35, http://dx.doi.org/10.1007/978-3-642-31104-8_3.
- [18] O. FEINERMAN, A. KORMAN. *Memory Lower Bounds for Randomized Collaborative Search and Implications for Biology*, in "Proceedings of the 26th International Symposium on Distributed Computing (DISC 2012)", 2012, p. 61–75, http://dx.doi.org/10.1007/978-3-642-33651-5_5.
- [19] O. FEINERMAN, A. KORMAN, Z. LOTKER, J. S. SERENI. *Collaborative search on the plane without communication*, in "Proceedings of the 31st ACM symposium on Principles of distributed computing (PODC 2012)", 2012, p. 77–86, <http://doi.acm.org/10.1145/2332432.2332444>.
- [20] P. FRAIGNIAUD, A. KORMAN, M. PARTER, D. PELEG. *Randomized Distributed Decision*, in "Proceedings of the 26th International Symposium on Distributed Computing (DISC 2012)", 2012, p. 371–385, http://dx.doi.org/10.1007/978-3-642-33651-5_26.
- [21] P. FRAIGNIAUD, A. PELC. *Decidability Classes for Mobile Agents Computing*, in "LATIN", 2012, p. 362–374.
- [22] M. HABIB, A. MAMCARZ, F. DE MONTGOLFIER. *Algorithms for Some H-Join Decompositions*, in "LATIN", 2012, p. 446–457.

National Conferences with Proceeding

- [23] F. DE MONTGOLFIER, M. SOTO, L. VIENNOT. *Modularité asymptotique de quelques classes de graphes*, in "14èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications (ALGOTEL)", La Grande Motte, 2012, p. 1–4.

Research Reports

- [24] F. BACCELLI, F. MATHIEU, I. NORROS. *Spatial Interactions of Peers and Performance of File Sharing Systems*, Inria, May 2012, n^o RR-7713, 17, Accepted to Infocom 2013, <http://hal.inria.fr/inria-00615523>.
- [25] F. DURAND, F. MATHIEU, L. NOIRIE. *On the Manipulability of Voting Systems: Application to Multi-Carrier Networks*, LINC, April 2012, n^o 2012-04-001, 17, <http://hal.inria.fr/hal-00692096>.
- [26] N. HEGDE, L. MASSOULIÉ, L. VIENNOT. *Self-Organizing Flows in Social Networks*, Inria, 2012, p. 1–17.