IN PARTNERSHIP WITH:
CNRS

Université Denis Diderot (Paris 7)

Activity Report 2013

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

Creation of the Project-Team: 2007 July 01.

1. Members

Research Scientists
Laurent Viennot [Team leader, Inria, Senior Researcher, HdR]
Dominique Fortin [Inria, Researcher]
Pierre Fraigniaud [CNRS, Senior Researcher, HdR]
Amos Korman [CNRS, Researcher]

Faculty Members
Yacine Boufkhad [Univ. Paris VII, Associate Professor]
Fabien de Montgolfier [Univ. Paris VII, Associate Professor]
Michel Habib [Univ. Paris VII, Professor, HdR]
Hugues Fauconnier [Univ. Paris VII, Associate Professor, HdR]
Carole Delporte-Gallet [Univ. Paris VII, Professor, HdR]

External Collaborators
Fabien Mathieu [Bell Labs (Alcatel), HdR]
Ider Tseveendorj [Univ. Versailles, Associate Professor, HdR]

Engineer
François Durand [Inria]

PhD Students
The-Dang Huynh [Bell Labs (Alcatel)]
Leonardo Linguaglossa [Inria, Alcatel-Lucent Bell Labs, from May 2013]

Administrative Assistant
Christine Anocq [Inria, TRE 20% shared with Hypercom and Aoste]

2. Overall Objectives

2.1. Overall Objectives

GANG focuses on algorithm design for large scale networks using structural properties of these networks. Application domains include the development of optimized protocols for large dynamic networks such as mobile networks or overlay networks over Internet. This includes for instance peer-to-peer applications, or the navigability of social networks. GANG tools come from recent advances in the field of graph algorithms, both in centralized and distributed settings. In particular, this includes graph decomposition and geometric properties (such as low doubling dimension, low dimension embedding, etc.). Today, the management of large networks, Internet being the reference, is best effort. However, the demand for mobility (ad hoc networks, wireless connectivity, etc.) and for dynamicity (node churn, fault tolerance, etc.) is increasing. In this distributed setting, it becomes necessary to design a new generation of algorithms and protocols to face the challenge of large scale mobility and dynamicity. In the mean time, recent and sophisticated theoretical results have emerged, offering interesting new tracks for managing large networks. These results concern centralized and decentralized algorithms for solving key problems in communication networks, including routing, but also information retrieval, localization, or load balancing. They are mainly based on structural properties observed in most of real networks: approximate topology with low dimension metric spaces, low treewidth, low doubling dimension, graph minor freeness, etc. In addition, graph decomposition techniques have recently progressed. The scientific community has now tools for optimizing network management. First striking results include designing overlay networks for peer-to-peer systems and understanding the navigability of large social networks.

3. Research Program

3.1. Research Program

Taking into account the scientific achievements of the last years, and the short presentation section above, GANG is currently focusing on the following objectives:

- Graphs algorithms
- Distributed Computing
- P2P-like Algorithms for Future Networks

3.1.1. Graph algorithms

3.1.1.1. Graph Decompositions

We study new decompositions schemes such as 2-join, skew partitions and others partition problems. These graph decompositions appeared in the structural graph theory and are the basis of some well-known theorems such as the Perfect Graph Theorem. For these decompositions there is a lack of efficient algorithms. We aim at designing algorithms working in $O(nm)$ since we think that this could be a lower bound for these decompositions.

3.1.1.2. Graph Search

We more deeply study multi-sweep graph searches. In this domain a graph search only yields a total ordering of the vertices which can be used by the subsequent graph searches. This technique can be used on huge graphs and do not need extra memory. We already have obtained preliminary results in this direction and many well-known graph algorithms can be put in this framework. The idea behind this approach is that each sweep discovers some structure of the graph. At the end of the process either we have found the underlying structure (for example an interval representation for an interval graph) or an approximation of it (for example in hard discrete optimization problems). Application to exact computations of centers in huge graphs, to underlied combinatorial optimization problems, but also to networks arising in Biology.

3.1.2. Distributed computing

The distributed community can be viewed as the union of two sub-communities. This is true even in our team. Even though they are not completely disjoint, they are disjoint enough not to leverage each other's results. At a high level, one is mostly interested in timing issues (clock drifts, link delays, crashes, etc.) while the other one is mostly interested in spatial issues (network structure, memory requirements, etc.). Indeed, one sub-community is mostly focusing on the combined impact of asynchronism and faults on distributed computation, while the other addresses the impact of network structural properties on distributed computation. Both communities address various forms of computational complexities, through the analysis of different concepts. This includes, e.g., failure detectors and wait-free hierarchy for the former community, and compact labeling schemes and computing with advice for the latter community. We have the ambitious project to achieve the reconciliation between the two communities by focusing on the same class of problems, the yes/no-problems, and establishing the scientific foundations for building up a consistent theory of computability and complexity for distributed computing. The main question addressed is therefore: is the absence of globally coherent computational complexity theories covering more than fragments of distributed computing, inherent to the field? One issue is obviously the types of problems located at the core of distributed computing. Tasks like consensus, leader election, and broadcasting are of very different nature. They are not *yes-no* problems, neither are they minimization problems. Coloring and Minimal Spanning Tree are optimization problems but we are often more interested in constructing an optimal solution than in verifying the correctness of a given solution. Still, it makes full sense to analyze the *yes-no* problems corresponding to checking the validity of the output of tasks. Another issue is the power of individual computation. The FLP impossibility result as well as Linial's lower bound hold independently from the individual computational power of the involved computing entities. For instance, the individual power of solving NP-hard problems in constant time would not help overcoming these limits which are inherent to the fact that computation is distributed. A third issue is the abundance of models for distributed computing frameworks, from shared memory to message passing, spanning all kinds of specific network structures (complete graphs, unit-disk graphs, etc.) and or timing constraints (from complete synchronism to full asynchronism). There are however models, typically the wait-free model and the LOCAL model, which, though they do not claim to reflect accurately real distributed computing systems, enable focusing on some core issues. Our research program is ongoing to carry many important notions of Distributed Computing into a *standard* computational complexity.

3.1.3. A Peer-to-Peer approach to future content Distribution

Unexpectedly, the field of P2P applications is still growing and challenging issues remain worth studying.

3.1.3.1. New network models

The new models that have been proposed to take into account the evolution of network architecture and usage indicate new opportunities for P2P, like the possibility to have superscalable systems whose performance increases with the popularity. This surprising property, if it can be enforced, will give P2P an additional asset compared to the current situation. However, this results are still at an early stage, and it is planned to continue the study from a theoretical point of view, but also with experimentations with emulation and/or simulation of future networks on large grids.

3.1.3.2. P2P storage

The challenges of a persistent and robust distributed storage with respect to failures are nowadays relatively well understood. However, the results about instant availability are still not completely understood: how to give guarantees, in a P2P system where peers are not online 100% of the time, that a content will be available when its owner asks for it? Can we propose some allocation policy that ensures maximal availability with only a partial knowledge of online patterns? We believe that these issues, halfway between failure tolerance and opportunistic networks, are still promising.

3.1.3.3. Caching allocation

Today, most of content distribution is ensured by so-called Content Distribution Networks (CDNs). It is expected that caching techniques will remain a hot topic in the years to come, for instance through the studies related to Content Centric Networking, which is inspired by P2P content distribution paradigms, like using so-called chunks as the basic data exchange unit. Many challenges in this field are related to dimensioning and caching strategies. In GANG, we aim at conducting a study centered on the trade-offs between storage and bandwidth usage. Note that many studies have been/are realized on this topic, mostly rely on operational research methodology and offer solutions that can sometimes be difficult to use in practice. GANG uses a different approach, based on alternate modeling assumptions inherited from our previous achievments on bandwidth dimensioning. The goal of this complementary approach is to provide simple dimensioning guidelines while giving approximated, yet meaningful, performance evaluation.

3.1.3.4. Long term perspective on P2P content distribution

The success of YouTube-like delivery platforms (YouTube, DailyMotion...) does not come only from their technical performances, but from the ergonomy: these platforms allow to launch a video directly from one's browser, without the usual burden that comes with traditional P2P applications (install a specific client, open incoming ports, find .torrent files . . .). It is therefore important to keep working on basic P2P research, especially as many challenges are still open (see above), and new opportunities are likely to rise. First, advances in other fields may make P2P more interesting than other solutions –again. For instance, CCN protocols are designed to facilitate data dissemination. One could hope they open the way to CCN-assisted P2P protocols, where both the issues of ergonomy and network burden would be taken care of by design.

Unpredictable events, such as emerging/closing centralized filesharing services, can also change the power balance very fast with effects that are still hard to determine. For all these reasons, GANG aims at improving its expertise in the field of decentralized content distribution, even if it is quite difficult at the fast evolving current time to tell if that expertise should apply on traditional P2P, CCN, Cloud... architectures, or on any hybridation of these.

4. Application Domains

4.1. 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 of such studies.

5. New Results

5.1. Understanding graph representations

5.1.1. Connected graph searching

5.1.1.1. Computing H-Joins with Application to 2-Modular Decomposition

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

We present in [10], a general framework to design algorithms that compute H-join. For a given bipartite graph H, we say that a graph G admits a H-join decomposition or simply a H-join, if the vertices of G can be partitioned in |H| parts connected as in H. This graph H is a kind of pattern, that we want to discover in G. This framework allows us to present fastest known algorithms for the computation of P 4-join (aka N-join), P 5-join (aka W-join), C 6-join (aka 6-join). We also generalize this method to find a homogeneous pair (also known as 2-module), a pair M 1,M 2 such that for every vertex x∉(M 1∪M 2) and i∈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 in Graphs Comb. 3:127-139, 1987), it is a generalization of splits (a.k.a. 1-joins) and of modules. The algorithmics to compute them appears quite involved. In this paper, we describe an O(mn 2)-time algorithm computing all maximal homogeneous pairs of a graph, which not only improves a previous bound of O(mn 3) for finding only one pair (Everett et al. in Discrete Appl. Math. 72:209-218, 1997), but also uses a nice structural property of homogenous pairs, allowing to compute a canonical decomposition tree for sesquiprime graphs (i.e., graphs G having no module and such that for every vertex v∈G, G-v also has no module).

5.1.1.2. Algorithmic Aspects of Switch Cographs

Participants: Vincent Cohen-Addad, Michel Habib, Fabien de Montgolfier.

The paper [27], introduces the notion of involution module, the first generalization of the modular decomposition of 2-structure which has a unique linear-sized decomposition tree. We derive an $O(n^2)$ decomposition algorithm and we take advantage of the involution modular decomposition tree to state several algorithmic results. Cographs are the graphs that are totally decomposable w.r.t modular decomposition. In a similar way, we introduce the class of switch cographs, the class of graphs that are totally decomposable w.r.t involution modular decomposition. This class generalizes the class of cographs and is exactly the class of (Bull, Gem, Co-Gem, C_5)-free graphs. We use our new decomposition tool to design three practical algorithms for the maximum cut, vertex cover and vertex separator problems. The complexity of these problems was still unknown for this class of graphs. This paper also improves the complexity of the maximum clique, the maximum independant set, the chromatic number and the maximum clique cover problems by giving efficient algorithms, thanks to the decomposition tree. Eventually, we show that this class of graphs has Clique-Width at most 4 and that a Clique-Width expression can be computed in linear time.

5.1.1.3. LDFS-Based Certifying Algorithm for the Minimum Path Cover Problem on Cocomparability Graphs

Participants: Derek Corneil, Dalton Barnaby, Michel Habib.

For graph $G(V, E)$, a minimum path cover (MPC) is a minimum cardinality set of vertex disjoint paths that cover V (i.e., every vertex of G is in exactly one path in the cover). This problem is a natural generalization of the Hamiltonian path problem. Cocomparability graphs (the complements of graphs that have an acyclic transitive orientation of their edge sets) are a well studied subfamily of perfect graphs that includes many popular families of graphs such as interval, permutation, and cographs. Furthermore, for every cocomparability graph G and acyclic transitive orientation of the edges of $\overline{G}$ there is a corresponding poset P_G; it is easy to see that an MPC of G is a linear extension of P_G that minimizes the bump number of P_G. Although there are directly graph-theoretical MPC algorithms (i.e., algorithms that do not rely on poset formulations) for various subfamilies of cocomparability graphs, notably interval graphs, until now all MPC algorithms for cocomparability graphs themselves have been based on the bump number algorithms for posets. In this paper [5], we present the first directly graph-theoretical MPC algorithm for cocomparability graphs; this algorithm is based on two consecutive graph searches followed by a certifying algorithm. Surprisingly, except for a lexicographic depth first search (LDFS) preprocessing step, this algorithm is identical to the corresponding algorithm for interval graphs. The running time of the algorithm is $O(\min(n^2, n + \text{mloglogn}))$, with the nonlinearity coming from LDFS.

5.1.1.4. Easy identification of generalized common and conserved nested intervals

Participants: Fabien de Montgolfier, Mathieu Raffinot, Irena Rusu.

In the paper [28], we explain how to easily compute gene clusters, formalized by classical or generalized nested common or conserved intervals, between a set of K genomes represented as K permutations. A $b-$nested common (resp. conserved) interval I of size $|I|$ is either an interval of size 1 or a common (resp. conserved) interval that contains another b-nested common (resp. conserved) interval of size at least $|I| - b$. When $b = 1$, this corresponds to the classical notion of nested interval. We exhibit two simple algorithms to output all b-nested common or conserved intervals between K permutations in $O(Kn + \text{nocc})$ time, where nocc is the total number of such intervals. We also explain how to count all b-nested intervals in O(Kn) time. New properties of the family of conserved intervals are proposed to do so.

5.1.1.5. On computing the diameter of real-world undirected graphs

Participants: Pierluigi Crescenzi, Roberto Grossi, Michel Habib, Leonardo Lanzi, Andrea Marino.

We propose in [2],a new algorithm for the classical problem of computing the diameter of undirected unweighted graphs, namely, the maximum distance among all the pairs of nodes, where the distance of a pair of nodes is the number of edges contained in the shortest path connecting these two nodes. Although its worst-case complexity is O(nm) time, where n is the number of nodes and m is the number of edges of the graph, we experimentally show that our algorithm works in O(m) time in practice, requiring few breadth-first searches to complete its task on almost 200 real-world graphs.

5.1.1.6. Toward more localized local algorithms: removing assumptions concerning global knowledge

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

Numerous sophisticated local algorithm were suggested in the literature for various fundamental problems. Notable examples are the MIS and $(\Delta + 1)$-coloring algorithms by Barenboim and Elkin, by Kuhn, and by Panconesi and Srinivasan, as well as the $o(\Delta^2)$-coloring algorithm by Linial. Unfortunately, most known local algorithms (including, in particular, the aforementioned 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. This paper [11], provides a rather general method for transforming a non-uniform local algorithm into a *uniform* one. Furthermore, the resulting algorithm enjoys the same asymptotic running time as the original non-uniform algorithm. Our method applies to a wide family of both deterministic and randomized algorithms. Specifically, it applies to almost all of the state of the art non-uniform algorithms regarding MIS and Maximal Matching, as well as to many results concerning the coloring problem. (In particular, it applies to all aforementioned algorithms.) To obtain our transformations we introduce a new distributed tool called *pruning* algorithms, which we believe may be of independent interest.

5.1.2. Self-organizing Flows in 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 this paper [24], 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.

5.2. Large Scale Networks Performance and Modeling

5.2.1. Can P2P Networks be Super-Scalable?

Participants: François Baccelli, Fabien Mathieu, Ilkka Norros, Rémi Varloot.

We propose in [14], a new model for peer-to-peer networking which takes the network bottlenecks into account beyond the access. This model can cope with key features of P2P networking like degree or locality constraints together with the fact that distant peers often have a smaller rate than nearby peers. Using a network model based on rate functions, we give a closed form expression of peers download performance in the system's fluid limit, as well as approximations for the other cases. Our results show the existence of realistic settings for which the average download time is a decreasing function of the load, a phenomenon that we call super-scalability.

5.2.2. Contenu généré par les utilisateurs : une étude sur DailyMotion

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

Actuellement, une large part du trafic Internet vient de sites de "User-Generated Content" (UGC). Comprendre les caractéristiques de ce trafic est important pour les opérateurs (dimensionnement réseau), les fournisseurs (garantie de la qualité de service) et les équipementiers (conception d'équipements adaptés). Dans ce contexte, nous proposons [15], d'analyser et de modéliser des traces d'usage du site DailyMotion.

5.2.3. *Rumor Spreading in Random Evolving Graphs*

Participants: Andrea Clementi, Pierluigi Crescenzi, Carola Doerr, Pierre Fraigniaud, Isopi Marco, Alessandro Panconesi, Pasquale Francesco, Silvestri Riccardo.

In [13], we aim at analyzing the classical information spreading "push" protocol in *dynamic* networks. We consider the *edge-Markovian* evolving graph model which captures natural temporal dependencies between the structure of the network at time t, and the one at time $t+1$. Precisely, a non-edge appears with probability p, while an existing edge dies with probability q. In order to fit with real-world traces, we mostly concentrate our study on the case where $p = \Omega(\frac{1}{n})$ and q is constant. We prove that, in this realistic scenario, the "push" protocol does perform well, completing information spreading in $O(\log n)$ time steps, w.h.p., even when the network is, w.h.p., disconnected at every time step (e.g., when $p \ll \frac{\log n}{n}$). The bound is tight. We also address other ranges of parameters p and q (e.g., $p+q=1$ with arbitrary p and q, and $p = \Theta\left(\frac{1}{n}\right)$ with arbitrary q). Although they do not precisely fit with the measures performed on real-world traces, they can be of independent interest for other settings. The results in these cases confirm the positive impact of dynamism.

5.3. Complexity issues in distributed graph algorithms

5.3.1. *What can be decided locally without identifiers?*

Participants: Pierre Fraigniaud, Mika Göös, Amos Korman, Jukka Suomela.

Do unique node identifiers help in deciding whether a network G has a prescribed property P? We study this question in the context of distributed local decision, where the objective is to decide whether $G \in P$ by having each node run a constant-time distributed decision algorithm. If $G \in P$, all the nodes should output yes; if $G \notin P$, at least one node should output no. A recent work (Fraigniaud et al., OPODIS 2012) studied the role of identifiers in local decision and gave several conditions under which identifiers are not needed. In this article [21], we answer their original question. More than that, we do so under all combinations of the following two critical variations on the underlying model of distributed computing: (B): the size of the identifiers is bounded by a function of the size of the input network; as opposed to ($\neg B$): the identifiers are unbounded. (C): the nodes run a computable algorithm; as opposed to ($\neg C$): the nodes can compute any, possibly uncomputable function. While it is easy to see that under ($\neg B, \neg C$) identifiers are not needed, we show that under all other combinations there are properties that can be decided locally if and only if identifiers are present. Our constructions use ideas from classical computability theory.

5.3.2. *Local Distributed Decision*

Participants: Pierre Fraigniaud, Amos Korman, David Peleg.

A central theme in distributed network algorithms concerns understanding and coping with the issue of locality. Inspired by sequential complexity theory, we focus on a complexity theory for distributed decision problems. In the context of locality, solving a decision problem requires the processors to independently inspect their local neighborhoods and then collectively decide whether a given global input instance belongs to some specified language. Our paper [7], introduces several classes of distributed decision problems, proves separation among them and presents some complete problems. More specifically, we consider the standard LOCAL model of computation and define LD (for local decision) as the class of decision problems that can be solved in constant number of communication rounds. We first study the intriguing question of whether randomization helps in local distributed computing, and to what extent. Specifically, we define the corresponding randomized class BPLD, and ask whether LD=BPLD. We provide a partial answer to this question by showing that in many cases, randomization does not help for deciding hereditary languages. In addition, we define the notion of local many-one reductions, and introduce the (nondeterministic) class NLD of decision problems for which there exists a certificate that can be verified in constant number of communication rounds. We prove that there exists an NLD-complete problem. We also show that there exist problems not in NLD. On the other hand, we prove that the class NLD#n, which is NLD assuming that each processor can access an oracle that provides the number of nodes in the network, contains all (decidable) languages. For this class we provide a natural complete problem as well.

5.3.3. Locality and checkability in wait-free computing

Participants: Pierre Fraigniaud, Sergio Rajsbaum, Travers Corentin.

The paper [9],studies notions of locality that are inherent to the specification of distributed tasks, and independent of the computing model, by identifying fundamental relationships between the various scales of computation, from the individual process to the whole system. A locality property called *projection-closed* is identified. This property completely characterizes tasks that are wait-free *checkable*, where a task $T = (\mathcal{I}, \mathcal{O}, \Delta)$ is said to be checkable if there exists a distributed algorithm that, given $s \in \mathcal{I}$ and $t \in \mathcal{O}$, determines whether $t \in \Delta(s)$, i.e., whether t is a valid output for s according to the specification of T. Projection-closed tasks are proved to form a rich class of tasks. In particular, determining whether a projection-closed task is wait-free solvable is shown to be undecidable. A stronger notion of locality is identified by considering tasks whose outputs "look identical" to the inputs at every process: a task $T = (\mathcal{I}, \mathcal{O}, \Delta)$ is said to be *locality-preserving* if $\mathcal{O}$ is a covering complex of $\mathcal{I}$. We show that this topological property yields obstacles for wait-free solvability different in nature from the classical impossibility results. On the other hand, locality-preserving tasks are projection-closed, and thus they are wait-free checkable. A classification of locality-preserving tasks in term of their relative computational power is provided. This is achieved by defining a correspondence between subgroups of the *edgepath* group of an input complex and locality-preserving tasks. This correspondence enables to demonstrate the existence of hierarchies of locality-preserving tasks, each one containing, at the top, the universal task (induced by the universal covering complex), and, at the bottom, the trivial identity task.

5.3.4. Delays Induce an Exponential Memory Gap for Rendezvous in Trees

Participants: Pierre Fraigniaud, Pelc Andrzej.

The aim of rendezvous in a graph is meeting of two mobile agents at some node of an unknown anonymous connected graph. In this paper [8], we focus on rendezvous in trees, and, analogously to the efforts that have been made for solving the exploration problem with compact automata, we study the size of memory of mobile agents that permits to solve the rendezvous problem deterministically. We assume that the agents are identical, and move in synchronous rounds. We first show that if the delay between the starting times of the agents is *arbitrary*, then the lower bound on memory required for rendezvous is $\Omega(\log n)$ bits, even for the line of length n. This lower bound meets a previously known upper bound of $O(\log n)$ bits for rendezvous in arbitrary graphs of size at most n. Our main result is a proof that the amount of memory needed for rendezvous *with simultaneous start* depends essentially on the number ℓ of leaves of the tree, and is exponentially less impacted by the number n of nodes. Indeed, we present two identical agents with $O(\log \ell + \log \log n)$ bits of memory that solve the rendezvous problem in all trees with at most n nodes and at most ℓ leaves. Hence, for the class of trees with polylogarithmically many leaves, there is an exponential gap in minimum memory size needed for rendezvous between the scenario with arbitrary delay and the scenario with delay zero. Moreover, we show that our upper bound is optimal by proving that $\Omega(\log \ell + \log \log n)$ bits of memory are required for rendezvous, even in the class of trees with degrees bounded by 3.

5.3.5. On the Manipulability of Voting Systems: Application to Multi-Operator Networks

Participants: François Durand, Fabien Mathieu, Ludovic Noirie.

Internet is a large-scale and highly competitive economic ecosystem. In order to make fair decisions, while preventing the economic actors from manipulating the natural outcome of the decision process, game theory is a natural framework, and voting systems represent an interesting alternative that, to our knowledge, has not yet being considered. They allow competing entities to decide among different options. In this paper [20], we investigate their use for end-to-end path selection in multi-operator networks, analyzing their manipulability by tactical voting and their economic efficiency.We show that Instant Runoff Voting is much more efficient and resistant to tactical voting than the natural system which tries to get the economic optimum.

5.4. Communication and Fault Tolerance in Distributed Networks

5.4.1. Linear Space Bootstrap Communication Schemes

Participants: Carole Delporte-Gallet, Hugues Fauconnier, Eli Gafni, Sergio Rajsbaum.

We consider in [18], a system of n processes with ids not a priori known, that are drawn from a large space, potentially unbounded. How can these n processes communicate to solve a task? We show that n a priori allocated Multi-Writer Multi-Reader (MWMR) registers are both needed and sufficient to solve any read-write wait free solvable task. This contrasts with the existing possible solution borrowed from adaptive algorithms that require $\Theta(n^2)$ MWMR registers. To obtain these results, the paper shows how the processes can non blocking emulate a system of n Single-Writer Multi-Reader (SWMR) registers on top of n MWMR registers. It is impossible to do such an emulation with $n-1$ MWMR registers. Furthermore, we want to solve a sequence of tasks (potentially infinite) that are sequentially dependent (processes need the previous task's outputs in order to proceed to the next task). A non blocking emulation might starve a process forever. By doubling the space complexity, using $2n-1$ rather than just n registers, the computation is wait free rather than non blocking.

5.4.2. Black Art: Obstruction-Free k-set Agreement with |MWMR registers| < |proccesses|

Participants: Carole Delporte-Gallet, Hugues Fauconnier, Eli Gafni, Sergio Rajsbaum.

When n processes communicate by writing to and reading from $k < n$ MWMR registers the "communication bandwidth" precludes emulation of SWMR system, even non-blocking.

Nevertheless, recently a positive result was shown that such a system either wait-free or obstruction-free can solve an interesting one-shot task. This paper demonstrates another such result. It shows that $(n-1)-$set agreement can be solved obstruction-free with merely 2 MWMR registers. Achieving $k-$set agreement with $n-k+1$ registers is a challenge. In [17], we make the first step toward it by showing $k-$set agreement with $2(n-k)$ registers.

5.4.3. Adaptive Register Allocation with a Linear Number of Registers

Participants: Carole Delporte-Gallet, Hugues Fauconnier, Eli Gafni, Leslie Lamport.

In [16], we give an adaptive algorithm in which processes use multi-writer multi- reader registers to acquire exclusive write access to their own single-writer, multi-reader registers. It is the first such algorithm that uses a number of registers linear in the number of participating processes. Previous adaptive algorithms require at least $\Theta(n^{3/2})$ registers

5.4.4. Uniform Consensus with Homonyms and Omission Failures

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

In synchronous message passing models in which some processes may be homonyms, i.e. may share the same id, we consider the consensus problem. Many results have already been proved concerning Byzantine failures in models with homonyms, we complete in [19], the picture with crash and omission failures.

Let n be the number of processes, t the number of processes that may be faulty ($t < n$) and l ($1 \leq l \leq n$) the number of identifiers. We prove that for crash failures and send-omission failures, uniform consensus is solvable even if $l = 1$, that is with fully anonymous processes for any number of faulty processes.

Concerning omission failures, when the processes are numerate, i.e. are able to count the number of copies of identical messages they received in each round, uniform consensus is solvable even for fully anonymous processes for $n > 2t$. If processes are not numerate, uniform consensus is solvable if and only if $l > 2t$.

All the proposed protocols are optimal both in the number of communication steps needed, and in the number of processes that can be faulty.

All these results show, (1) that identifiers are not useful for crash and send-omission failures or when processes are numerate, (2) for general omission or for Byzantine failures the number of different ids becomes significant.

5.4.5. Byzantine agreement with homonyms

Participants: Carole Delporte-Gallet, Hugues Fauconnier, Rachid Guerraoui, Anne-Marie Kermarrec, Hung Tran-The.

So far, the distributed computing community has either as- sumed 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 dif- ferent authenticated identifiers, where $1 \leq l \leq n$. In this paper [3], 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 sur- prisingly, the number of identifiers must be greater than $(n + 3t)/2$ in the partially synchronous case. This demonstrates two differences from the classical model (which has l = n): there are situations where relaxing synchrony to partial syn- chrony renders agreement impossible; and, in the partially synchronous case, increasing the number of correct processes can actually make it harder to reach agreement. The im- possibility 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.

5.4.6. Byzantine agreement with homonyms in synchronous systems

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

We consider in [4], the Byzantine agreement problem 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 $n_1 ..., n_l$ 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 $l\frac{(n-r)t}{n-t-\min(t,r)}$ where $r = n \mod l$.

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.

5.4.7. Convergence of the D-iteration algorithm: convergence rate and asynchronous distributed scheme

Participants: Dohy Hong, Fabien Mathieu, Gérard Burnside.

In this paper [25], we define the general framework to describe the diffusion operators associated to a positive matrix. We define the equations associated to diffusion operators and present some general properties of their state vectors. We show how this can be applied to prove and improve the convergence of a fixed point problem associated to the matrix iteration scheme, including for distributed computation framework. The approach can be understood as a decomposition of the matrix-vector product operation in elementary operations at the vector entry level.

5.5. Discrete Optimization Algorithms

5.5.1. Shrinking Maxima, Decreasing Costs: New Online Packing and Covering Problems

Participants: Pierre Fraigniaud, Magnús M. Halldórsson, Boaz Patt-Shamir, Dror Rawitz, Adi Rosén.

We consider in [23], two new variants of online integer programs that are duals. In the packing problem we are given a set of items and a collection of knapsack constraints over these items that are revealed over time in an online fashion. Upon arrival of a constraint we may need to remove several items (irrevocably) so as to maintain feasibility of the solution. Hence, the set of packed items becomes smaller over time. The goal is to maximize the number, or value, of packed items. The problem originates from a buffer-overflow model in communication networks, where items represent information units broken into multiple packets. The other problem considered is online covering: There is a universe to be covered. Sets arrive online, and we must decide for each set whether we add it to the cover or give it up. The cost of a solution is the total cost of sets taken, plus a penalty for each uncovered element. The number of sets in the solution grows over time, but its cost goes down. This problem is motivated by team formation, where the universe consists of skills, and sets represent candidates we may hire. The packing problem was introduced for the special case where the matrix is binary; in this paper we extend the solution to general matrices with non-negative integer entries. The covering problem is introduced in this paper; we present matching upper and lower bounds on its competitive ratio.

5.5.2. Generalized Subdifferentials of the Sign Change Counting Function

Participants: Dominique Fortin, Ider Tseveendorj.

A natural generalization of piecewise linear approximation of non convex problems relies on piecewise convex approximation; along the way to solve the piecewise convex maximization problem [30] both effectively and efficiently, optimality conditions have to be addressed in two ways: either the violation of necessary conditions should lead to a direction of improvement from a local solution, or a sufficient condition for global optimality has to be fulfilled. The way to either goal is paved with subdifferentials and their generalizations on a per problem basis.

In the article [29], the counting function on binary values is extended to the signed case in order to count the number of transitions between contiguous locations. A generalized subdifferential for the sign change counting function is given where classical subdifferentials remain intractable. An attempt to prove global optimality at some point, for the 4-dimensional first non trivial example, is made by using a sufficient condition specially tailored among all the cases for this subdifferential.

6. Bilateral Contracts and Grants with Industry

6.1. Bilateral Contracts with Industry

6.1.1. Radiocéros

Participant: Fabien Mathieu.

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

6.1.2. Alcatel

Participants: François Durand, The-Dang Huynh, Leonardo Linguaglossa, Fabien Mathieu, Laurent Viennot.

Gang has a strong collaboration with Alcatel-Lucent. Fabien Mathieu has moved from Gang to Alcatel-Lucent in May 2013. We focus on three aspects of networks :

- François Durand is funded through an ADR with the LINCS for studying voting systems and how they can be used to take distributed decision in multipartite networks.
- The-Dang Huynh is funded through a CIFRE PhD. for developping pagerank techniques in the context of social networks.
- Leonardo Linguaglossa is funded through an ADR with Inria in the context of the joint laboratory for studying the feasibility of information centric networking with a special focus on routing aspects.

7. Partnerships and Cooperations

7.1. National Initiatives

7.1.1. 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.

7.1.2. 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.

7.1.3. Alcatel-Lucent Bell Labs and Inria Joint Research Lab

Participants: The-Dang Huynh, Leonardo Linguaglossa, Fabien Mathieu, Laurent Viennot.

Gang is participating to the joint laboratory between Alcatel-Lucent and Inria and contributes mainly in the ADR (joint research action) on content centric networking.

7.1.4. Laboratory of Information, Networking and Communication Sciences (LINCS)

Participants: The-Dang Huynh, Leonardo Linguaglossa, Fabien Mathieu, Laurent Viennot.

Gang is participating to the LINCS, a research centre co-founded by Inria, Institut Mines-Télécom, UPMC and Alcatel-Lucent Bell Labs, dedicated to research and innovation in the domains of future information and communication networks, systems and services. Most of the collaboration with Alcatel-Lucent is carried through this structure.

7.2. European Initiatives

7.2.1. FP7 Projects

7.2.1.1. 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 PolitÃ¨cnica 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.

7.3. International Initiatives

7.3.1. 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).

7.3.2. Inria International Partners

7.3.2.1. Informal International Partners

Participants: Carole Delporte, Hugues Fauconnier.

- distributed computing and synchronization: regular visits by Sam Toueg (Toronto), Rachid Guerraoui (EPFL) and Luis Rodriguez (U. Lisboa).
- consensus agreement: Last year we have shown that $(n-1)$-set consensus can be solved obstruction-free with 2 MWMR registers and this bound is tight. We have tried to generalize this result to the $(n-k)$-set consensus with $k+1$ registers; our regular cooperation with Eli Gafni (UCLA) is still ongoing.

8. Dissemination

8.1. Scientific Animation

Participants: Carole Delporte, Hugues Fauconnier.

Moth participants have been invited for the TRANSFORM Summer School on research Directions in Distributed Computing Heraklion, Greece 10-14 june http://www.ics.forth.gr/carv/transform/srdc/index.html

8.2. Teaching - Supervision - Juries

8.2.1. 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.

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, 33 hours;
- H. Fauconnier, Internet Protocols and Distributed algorithms, 44 hours;

Master Bioinformatique University of Paris Diderot:

- L. Viennot, System, network and Internet, 15 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;

Elementary school: L. Viennot, “Ateliers de fondements de l’informatique", 30 hours (CM1, corresponding to 4th grade).

8.2.2. Supervision

PhD :Hung Tran-The, "Le Problème du Consensus dans les Systèmes avec Homonymes", University of Paris Diderot June 06 2013, supervised by Hugues Fauconnier and Carole Delporte [1],

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, Dohy Hong and Laurent Viennot,
Leonardo Linguaglossa, "Design of algorithms and protocols to support ICN functionalities in high speed routers", since 2013, supervised by Fabien Mathieu, Diego Perino and Laurent Viennot.

8.2.3. Juries

HdR reviews: C. Delporte-Gallet reviews Coelho’s HdR thesis "Contributions à la performance du calcul scientique et embarqué" (october 11th 2013)
L. Viennot has reviewed Nicolas Bonichon’s HdR thesis "Quelques algorithmes entre le monde des graphes et les nuages de points" (April 3rd 2013) and Adrian Kosowski’s HdR thesis "Time and space-efficient algorithms for mobile agents in an anonymous network" (October 26th 2013).

PhD review: L. Viennot has reviewed Przemyslaw Uznanski’s PhD thesis "Large scale platform : instantiable models and algorithmic design of communication schemes" (October 11th 2013), Remigiusz Modrzejewski’s PhD thesis "Distribution and storage in networks" (October 24th 2013) and Christian Glacet’s PhD thesis "Algorithmes de routage, de la réduction des coûts de communication à la dynamique" (December 6th 2013).

8.3. Popularization

8.3.1. scientific popularization

Participant: Laurent Viennot.

Laurent Viennot is correspondent for the Inria Paris-Rocquencourt center concerning scientific popularization. In that context, he has participated to a document describing how scientific popularization actions could take place at Inria [26]

9. Bibliography

Publications of the year

Doctoral Dissertations and Habilitation Theses

[1] H. TRAN-THE. , *Problème du Consensus dans le Modèle Homonyme*, Université Paris-Diderot - Paris VII, June 2013, http://hal.inria.fr/tel-00925941

Articles in International Peer-Reviewed Journals

[2] P. CRESCENZI, R. GROSSI, H. MICHEL, L. LANZI, A. MARINO. *On computing the diameter of real-world undirected graphs*, in "Theoretical Computer Science", 2013, vol. 514, pp. 84-95 [*DOI :* 10.1016/J.TCS.2012.09.018], http://hal.inria.fr/hal-00936304

[3] C. DELPORTE-GALLET, H. FAUCONNIER, R. GUERRAOUI, A.-M. KERMARREC, E. RUPPERT, H. TRAN-THE. *Byzantine agreement with homonyms*, in "Distributed Computing", 2013, vol. 26, n^o 5-6, pp. 321-340 [*DOI :* 10.1007/S00446-013-0190-3], http://hal.inria.fr/hal-00839625

[4] C. DELPORTE-GALLET, H. FAUCONNIER, H. TRAN-THE. *Byzantine agreement with homonyms in synchronous systems*, in "Theoretical Computer Science", 2013, vol. 496, pp. 34-49 [*DOI :* 10.1016/J.TCS.2012.11.012], http://hal.inria.fr/hal-00922415

[5] C. DEREK, B. DALTON, H. MICHEL. *LDFS-Based Certifying Algorithm for the Minimum Path Cover Problem on Cocomparability Graphs*, in "SIAM Journal on Computing", 2013, vol. 42, n^o 3, pp. 792-807 [*DOI :* 10.1137/11083856X], http://hal.inria.fr/hal-00936300

[6] P. FRAIGNIAUD. *Special issue with selected papers from PODC 2011*, in "Distributed Computing", 2013, vol. 26, n^o 5-6, 271 p. , http://hal.inria.fr/hal-00922692

[7] P. FRAIGNIAUD, A. KORMAN, D. PELEG. *Local Distributed Decision*, in "Journal of the ACM", 2013, 35 p. , http://hal.inria.fr/hal-00912561

[8] P. FRAIGNIAUD, A. PELC. *Delays Induce an Exponential Memory Gap for Rendezvous in Trees*, in "ACM Transactions on Algorithms", 2013, vol. 9, n^o 2, 17 p. , http://hal.inria.fr/hal-00922693

[9] P. FRAIGNIAUD, S. RAJSBAUM, C. TRAVERS. *Locality and checkability in wait-free computing*, in "Distributed Computing", 2013, vol. 26, n^o 4, pp. 223-242, http://hal.inria.fr/hal-00922691

[10] M. HABIB, A. MAMCARZ, F. DE MONTGOLFIER. *Computing H-Joins with Application to 2-Modular Decomposition*, in "Algorithmica", September 2013, 22 p. [*DOI :* 10.1007/S00453-013-9820-1], http://hal.inria.fr/hal-00921775

[11] A. KORMAN, J.-S. SERENI, L. VIENNOT. *Toward more localized local algorithms: removing assumptions concerning global knowledge*, in "Distributed Computing", 2013, vol. 26, n^o 5-6, pp. 289-308, http://hal.inria.fr/hal-00909713

[12] A. KORMAN, K. SHAY. *Controller and estimator for dynamic networks*, in "Information and Computation", 2013, vol. 223, pp. 43-66 [*DOI :* 10.1016/J.IC.2012.10.018], http://hal.inria.fr/hal-00912550

International Conferences with Proceedings

[13] E. ANDREA, P. CRESCENZI, C. DOERR, P. FRAIGNIAUD, M. ISOPI, A. PANCONESI, F. PASQUALE, R. SILVESTRI. *Rumor Spreading in Random Evolving Graphs*, in "ESA", France, 2013, pp. 325-336, http://hal.inria.fr/hal-00922696

[14] F. BACCELLI, F. MATHIEU, I. NORROS, R. VARLOOT. *Can P2P Networks be Super-Scalable?*, in "IEEE Infocom 2013 - 32nd IEEE International Conference on Computer Communications", Turin, Italy, 2013, http://hal.inria.fr/hal-00817069

[15] Y. CARLINET, T. D. HUYNH, B. KAUFFMANN, F. MATHIEU, L. NOIRIE, S. TIXEUIL. *Contenu généré par les utilisateurs : une étude sur DailyMotion*, in "15èmes Rencontres Francophones sur les Aspects Algorithmiques des Télécommunications (AlgoTel)", Pornic, France, N. NISSE, F. ROUSSEAU, Y. BUSNEL (editors), 2013, pp. 1-4, http://hal.inria.fr/hal-00818606

[16] C. Delporte-Gallet, H. Fauconnier, E. Gafni, L. Lamport. *Adaptive Register Allocation with a Linear Number of Registers*, in "DISC 2013 - 27th International Symposium Distributed Computing", Jerusalem, Israel, Y. Afek (editor), Lecture Notes in Computer Science, Springer, 2013, vol. 8205, pp. 269-283 [*DOI :* 10.1007/978-3-642-41527-2_19], http://hal.inria.fr/hal-00922425

[17] C. Delporte-Gallet, H. Fauconnier, E. Gafni, S. Rajsbaum. *Black Art: Obstruction-Free k-set Agreement with |MWMR registers| < |proccesses|*, in "NETYS 2013 - First International Conference Networked Systems", Marrakech, Morocco, V. Gramoli, R. Guerraoui (editors), Lecture Notes in Computer Science, Springer, 2013, vol. 7853, pp. 28-41 [*DOI :* 10.1007/978-3-642-40148-0_3], http://hal.inria.fr/hal-00922423

[18] C. Delporte-Gallet, H. Fauconnier, E. Gafni, S. Rajsbaum. *Linear Space Bootstrap Communication Schemes*, in "ICDCN 2013 - 14th International Conference Distributed Computing and Networking", Mumbi, India, D. Frey, M. Raynal, S. Sarkar, R. K. Shyamasundar, P. Sinha (editors), Lecture Notes in Computer Science, Springer, 2013, vol. 7730, pp. 363-377 [*DOI :* 10.1007/978-3-642-35668-1_25], http://hal.inria.fr/hal-00922420

[19] C. Delporte-Gallet, H. Fauconnier, H. Tran-The. *Uniform Consensus with Homonyms and Omission Failures*, in "ICDCN 2013 - 14th International Conference Distributed Computing and Networking", Mumbai, India, D. Frey, M. Raynal, S. Sarkar, R. K. Shyamasundar, P. Sinha (editors), Lecture Notes in Computer Science, Springer, 2013, vol. 7730, pp. 161-175 [*DOI :* 10.1007/978-3-642-35668-1_12], http://hal.inria.fr/hal-00922428

[20] F. Durand, F. Mathieu, L. Noirie. *On the Manipulability of Voting Systems: Application to Multi-Operator Networks*, in "8th International Workshop on Internet Charging and QoS Technologies (ICQT 2013, collocated with the 9th International Conference on Network and Service Management CNSM 2013)", Zurich, Switzerland, October 2013, pp. 292-297, http://hal.inria.fr/hal-00874096

[21] P. Fraigniaud, M. Göös, A. Korman, J. Suomela. *What can be decided locally without identifiers?*, in "PODC'13 - ACM Symposium on Principles of Distributed Computing", Montreal, Canada, ACM, 2013, pp. 157-165 [*DOI :* 10.1145/2484239.2484264], http://hal.inria.fr/hal-00912527

[22] P. Fraigniaud, M. Göös, A. Korman, J. Suomela. *What can be decided locally without identifiers?*, in "PODC", France, 2013, pp. 157-165, http://hal.inria.fr/hal-00922698

[23] P. Fraigniaud, M. Magnús, B. Patt-Shamir, D. Rawitz, A. Rosén. *Shrinking Maxima, Decreasing Costs: New Online Packing and Covering Problems*, in "APPROX-RANDOM", France, 2013, pp. 158-172, http://hal.inria.fr/hal-00922695

[24] N. Hegde, L. Massoulié, L. Viennot. *Self-organizing Flows in Social Networks*, in "Structural Information AND Communication Complexity - 20th International Colloquium, SIROCCO", Ischia, Italy, Springer, 2013, vol. 8179, pp. 116-128, http://hal.inria.fr/hal-00909718

Research Reports

[25] D. Hong, F. Mathieu, G. Burnside. , *Convergence of the D-iteration algorithm: convergence rate and asynchronous distributed scheme*, January 2013, 9 p. , http://hal.inria.fr/hal-00776084

[26] A. ROUSSEAU, A. DARNAUD, B. GOGLIN, C. ACHARIAN, C. LEININGER, C. GODIN, C. HOLIK, C. KIRCHNER, D. RIVES, E. DARQUIE, E. KERRIEN, F. NEYRET, F. MASSEGLIA, F. DUFOUR, G. BERRY, G. DOWEK, H. ROBAK, H. XYPAS, I. ILLINA, I. GNAEDIG, J. JONGWANE, J. EHREL, L. VIENNOT, L. GUION, L. CALDERAN, L. KOVACIC, M. COLLIN, M.-A. ENARD, M.-H. COMTE, M. QUINSON, M. OLIVI, M. GIRAUD, M. DORÉMUS, M. OGOUCHI, M. DROIN, N. LACAUX, N. ROUGIER, N. ROUSSEL, P. GUITTON, P. PETERLONGO, R.-M. CORNUS, S. VANDERMEERSCH, S. MAHEO, S. LEFEBVRE, S. BOLDO, T. VIÉVILLE, V. POIREL, A. CHABREUIL, A. FISCHER, C. FARGE, C. VADEL, I. ASTIC, J.-P. DUMONT, L. FÉJOZ, P. RAMBERT, P. PARADINAS, S. DE QUATREBARBES, S. LAURENT. , *Médiation Scientifique : une facette de nos métiers de la recherche*, March 2013, 34 p. , http://hal.inria.fr/hal-00804915

Other Publications

[27] V. COHEN-ADDAD, M. HABIB, F. DE MONTGOLFIER. , *Algorithmic Aspects of Switch Cographs*, 2013, http://hal.inria.fr/hal-00921760

[28] F. DE MONTGOLFIER, M. RAFFINOT, I. RUSU. , *Easy identification of generalized common and conserved nested intervals*, 2013, http://hal.inria.fr/hal-00921762

[29] D. FORTIN, I. TSEVEENDORJ. , *Generalized Subdifferentials of the Sign Change Counting Function*, 2013, http://hal.inria.fr/hal-00915606

References in notes

[30] D. FORTIN, I. TSEVEENDORJ. *Attractive force search algorithm for piecewise convex maximization problems*, in "Optim. Lett.", 2012, vol. 6, n^{o} 7, pp. 1317–1333, http://dx.doi.org/10.1007/s11590-011-0395-y