INRIA, Evaluation of Theme Réseaux et Télécomunications

Project-team GANG

March 2012

Project-team title: GANG (Networks, Graphs and Algorithms)

Scientific leader: Laurent Viennot

Research center: Paris - Rocquencourt

Common project-team with: Paris Diderot University and CNRS

1 Personnel

Personnel (July 2007)

	Misc.	INRIA	CNRS	University	**Total**
DR (1) / Professors			1	1	**2**
CR (2) / Assistant Professors		2	1	3	**6**
Permanent Engineers (3)					
Temporary Engineers (4)					
PhD Students	4		1		**5**
Post-Doc.		1			**1**
Total	**4**	**3**	**3**	**4**	**14**
External Collaborators					
Visitors (> 1 month)					

(1) "Senior Research Scientist (Directeur de Recherche)"
(2) "Junior Research Scientist (Chargé de Recherche)"
(3) "Civil servant (CNRS, INRIA, ...)"
(4) "Associated with a contract (Ingénieur Expert or Ingénieur Associé)"

Personnel (March 2012)

	Misc.	INRIA	CNRS	University	**Total**
DR / Professors	1	1	2	2	**6**
CR / Assistant Professor		1		4	**5**
Permanent Engineer					
Temporary Engineer					
PhD Students	6				**6**
Post-Doc.					
Total	**7**	**2**	**2**	**6**	**17**
External Collaborators					
Visitors (> 1 month)					

Changes in staff

DR / Professors CR / Assistant Professors	Misc.	INRIA	CNRS	University	total
Arrival	1		1	2	5
Leaving			1		1

Comments: Emmanuelle Lebhar left the EPI in 2009 to go to Chili (currently back in France, but still on sabbatical). Amos Korman joined the project in 2008. Carole Delporte, Hugues Fauconnier and Fabien Mathieu joined the project in 2010.

Current composition of the project-team (March 2012):

- Laurent Viennot [Team leader, Senior Researcher DR Inria, HdR]
- Yacine Boufkhad [Assistant professor (Paris Diderot Univ.)]
- Pierre Charbit [Assistant professor (Paris Diderot Univ.)]
- Carole Delporte-Gallet [Professor (Paris Diderot Univ.), HdR]
- Dominique Fortin [Senior Researcher CR Inria]
- Hugues Fauconnier [Assistant Professor (Paris Diderot Univ.), HdR]
- Pierre Fraigniaud [Research Director (DR Cnrs), HdR]
- Michel Habib [Professor (Paris Diderot Univ.), HdR]
- Amos Korman [Senior Researcher Cnrs]
- Fabien Mathieu [Senior Researcher Inria, HdR]
- Fabien de Montgolfier [Assistant professor (Paris Diderot Univ.)]
- Heger Arfaoui [PhD Student]
- Hervé Baumann [PhD Student]
- Jérémie Dusart [PhD Student]
- Xavier Koegler [PhD Student]
- Antoine Mamcarz [PhD Student]
- Hung Tran-The [PhD Student]
- Christine Anocq [TRE shared time 20% (with Hipercom and Aoste)]

Current position of former project-team members (including PhD students during the 2008-2011 period):

- Julien Clément is quantitative analyst at Sophis
- Anh-Tuan Gai is CEO at CleverScale in Paris
- George Giakkoupis is a INRIA researcher (CR) at the ASAP EPI

- Thomas Hugel is a post-doc at Polytech'Nice and I3S (France)
- Emmanuelle Lebhar is currently teacher in Paris (France)
- Vincent Limouzy is assistant professor at Clermont-Ferrand (France)
- Diego Perino is Research Engineer at Alcatel-Lucent Bell Labs France
- Hoang-Anh Phan is working in the IT consulting business in Paris
- Mauricio Soto is ATER at Orléans (France)
- Thu-Hien To is post-doc at LIRMM (Montpellier, France)(France)

Last INRIA enlistments

- Fabien Mathieu, 2010 (Researcher)

Other comments:

2 Work progress

2.1 Keywords

Networks; graphs; algorithms

2.2 Context and overall goal of the project

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.

2.3 Objectives for the evaluation period

GANG objective is to push forward the overall tracks to extend recent theoretical results in the above domains. A long term goal is to enable distributed algorithms dedicated for practical applications such as:

- Internet management through metric models for latencies,
- new peer-to-peer applications such as cooperative streaming or peer-to-peer backup,
- managing large dynamic networks such as ad hoc networks or overlays over Internet.

In order to handle these objectives, the following scientific challenges were identified:

- Understanding Inherent Graph and Combinatorial Structure
- Tackling Large Scale Algorithmic Constraints
- Implementing Efficient Graph Structures

2.3.1 Evolution of initial objectives

During the last four years, some adjustments were made with respect to the initial objectives.

1. Graph algorithms has grown to become an objective of its own. Indeed, graph expertise is one of the main competitive advantages of GANG, compared to the other EPIs in the field of networks and telecommunications. This is for instance why INRIA direction asked us to start a collaboration with biologists interested in graph-related aspects of their work. Overall, it seemed important for us no to remain on our past laurels but to continue to address theoretical graph challenges.

2. The study of algorithms for large scale networks was identified as a whole research objective, covering (fully or partially) the initial fields *Internet management through metrics models*, *large dynamic networks*, *large scale algorithmic constraints.*

3. The initial objective related to peer-to-peer (P2P) systems was mainly focused on the structure of P2P overlays, on the side of indexing (Distributed Hash Tables) and on the side of unstructured, self-organized, systems. A shift appeared during the previous evaluation period towards the study of decentralized content distributed systems, motivated by: faster than expected progress with respect to the initial P2P objectives; strong interactions with industrial partners (France Telecom, Technicolor) interested by decentralized content distribution.

4. The arrival of Carole Delporte and Hugues Fauconnier has strengthened the expertise of GANG in distributed computing, creating a group dynamic within the EPI for the creation of a specific *distributed computing* objective.

It should be noted that although objectives (1) and (2) can share the same broad description (they both consider graph algorithms for large networks), they are very distinct in practice. The differences will be detailed in the following, but roughly we could say that the compass of objective (1) points to graphs, while the compass of objective (2) points to networks.

2.4 Objective (1): Graph algorithms

2.4.1 Personnel

Michel Habib, Yacine Boufkhad, Dominique Fortin, Fabien de Montgolfier, Vincent Limouzy, Thu-Hien To, Antoine Mamcarz, Jérémie Dusart.

2.4.2 Project-team positioning

Many research group working on graph algorithms focus on complexity results. Here we deal with the design of efficient graph algorithms that can be applied even on huge graphs.

Our tools are mainly graph searches which can be used to solve a problem or to preprocess the data when the optimization problem we deal with is NP-hard to solve.

New areas of applications recently appeared such as social networks, and complex nextworks for example in Biology and we are involved in these two areas (2012 ANR Algopol with sociologists and a long term collaboration with biologists on complex networks).

2.4.3 Scientific achievements

Graph decompositions Our group has a long and fruitful history on graph decomposition algorithms like for modular decomposition [43], and other related problems such as interval graph recognition, cograph recognition ...

During the evaluation period we achieve two important goals : a simple linear time algorithm for modular decomposition (ICALP 2008) and a linear time algorithm for split decomposition (to appear in SIAM J. of Discrete Math.).

δ-hyperbolicity We studied δ-hyperbolicity as an alternative polynomial parameter to capture the metrics of some networks. This nice parameter introduced by M. Gromov for groups can be computed in $O(|V|^4)$ and is both local and global. Moreover it can be applied with any graph metrics (i.e. on the usual discrete metrics, but also on any metrics obtained with some embedding of the graph on a surface). This leads to nice results on graph labellings and on polygons[13].

Graph traversal algorithms Following Corneil's seminal work on graph searches, we gave very elegant implementation of lexicographic depth first search (LexDFS) lexicographic breadth first search (LexBFS) through vertex partitioning.

Recently, new recognition algorithms have been devised using only LexBFS; it yields, in no more than 5 sweeps, for interval recognition and 2 sweeps for cograph recognition. Those 2 labelings (LexBFS and LexDFS) have been extended to different traversals (lexUp and LexDown) that could be implemented by the same vertex partitioning technique. It opens up a broad and promising field for tackling old and new special cases of optimization problems for the very near future, listed below. More generally, the graph traversal paradigm allows to go back to the more formal point of view of posets that are ubiquitous in the domain splitting approach to hard optimization problems that will be surveyed later on.

One of our best achievement is diameter computation on very large graphs. This exact algorithm starts with a *4-sweep heuristic*, which seems to be very efficient first approximation, and then produces in most of the graphs we consider the exact value of the diameter (with termination) after a few additional sweeps. For instance, the diameter of the Facebook network has been computed by Backstrom, Boldi, Rosa, Ugander, and Vigna by making use of the iFUB (iterative Fringe Upper Bound) exact method that we developed [62]. The diameter of Facebook is 41, and it was computed using only 17 breadth first searches!

Consecutive ones is a property that is satisfied for a binary matrix if there exists a column permutation such that the ones in each row are consecutive. It has a long history in biology applications (see for instance references in [94, 95]). It is related to the recognition of Robinson property in classification field up to a permutation, using so-called $PQ-$tree structure. Several algorithms with complexity $O(n^4)$, $O(n^3)$, and $O(n^2 \log n)$ for recognizing if a dissimilarity on n points is Robinsonian have been proposed. On the other hand, it is well-known that the tree metrics and the ultrametrics cannot be recognized in less than $O(n^2)$ time (which is proportional to the size of the input). In 2007, Pascal Préa (LIF Marseille) found a way to yield an optimal $O(n^2)$ recognition algorithm for Robinsonian dissimilarities, which led to a collaboration with Dominique Fortin.

The algorithm is based on the intimate (known) relationships between Robinsonian dissimilarities and interval hypergraphs and the data structure of $PQ-$trees used for recognizing interval graphs in linear time. We are currently considering zn alternate approach using LexBFS-like sweep algorithms, which seem to be able to compete with the current interval hypergraph framework. This would provide a linear programming modeling for the problem (to our knowledge, there exists only a piecewise convex formulation). Recently, we have started to extend the consecutive ones property to consecutive signs property $\{0, \pm 1\}$ entries instead. We are also considering the consecutive ones property formulated as a quadratic assignment problem (QAP). QAP is known as one of the hardest combinatorial problems (encompassing traveling salesman), similar to the maximum clique problem (see piecewise convex formulations below); for this new challenging property, first attempts in term of 2 QAPs fail to capture the actual difficulty and justify this application for long term research.

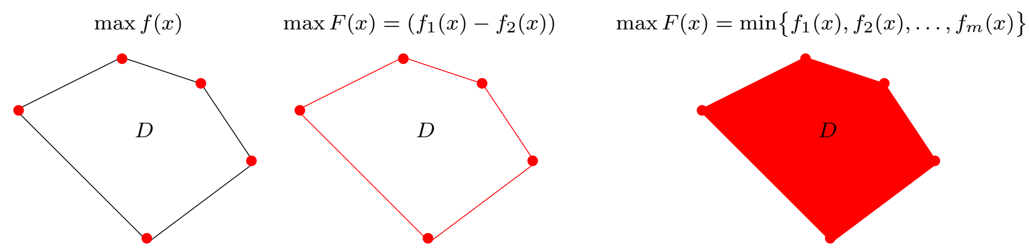

Figure 1: location of optimal solutions to various (non convex) problems involving convex functions over a polyhedron

Piecewise Convex Maximization Problems (structural properties) In the constant effort to solve non convex optimization problems, different classes have been addressed: convex maximization, difference of convex functions (DC) maximization and piecewise convex maximization (PCMP). These three classes present an increasing difficulty, as the location of the optimal solutions can be on the vertices, the boundaries or the whole domain respectively (cf Figure 1).

We focused on the PCMP class, because for some scenarios, it can outperform the other two classes (most problems can be modeled in any of the three formulations) [28, 34].

Given a local solution, all worse points have to be included in the domain to prove global optimality (see Figure 2). Since 2001, when Ider Tseveendorj gave the necessary and sufficient conditions for global optimality of PCMP, he and Dominique Fortin have constantly improved algorithms to solve the problem (hard even in low dimension), starting from tools from linear programming, then adding convex obstruction (analog to hyperplane cuts) [32, 35].

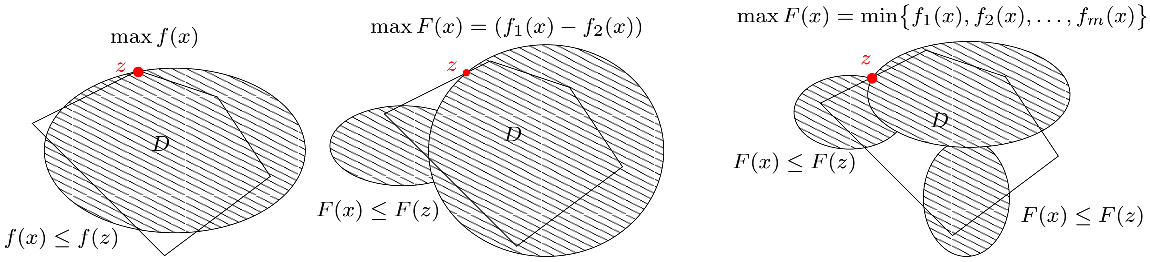

Figure 2: escaping from a local solution vs inclusion of worse points into the domain

Grothendieck's constant has stimulated a lot of research since its introduction. We are interested in it because of its applications for high dimensional convex maximization problem. In 1994, Jim Reeds (joint work with P.C. Fishburn) gave an upper bound of the lower bound by considering a convex maximization problem. During its internship S. Singh from IIT Kanpur worked with D. Fortin, on a family of Kronecker power of Hadamard matrices, generalizing the original idea by J. Reeds. The estimation of the constant was not improved, but the study allowed to address more families under the same structural orbit [25]. Jim Reeds (private communication) gave to us a different orbit that outperforms ours, for small sizes, but with the same asymptotical result.

Complex networks in Biology Thanks to A. Petit who gave us the contact, we have started a long term collaboration with a group of young biologists working at University Paris 6.

We first studied graph-based phylogeny, trying to generalize results on-tree based phylogeny to arbitrary graph structures. This subject was studied in T-H To's PhD. It leads to a proof of an old (20 years) conjecture from M. Steel [8]. Recently we considered clustering algorithms for bipartite graphs. The clusters has computed by our algorithm seem to have important biological meanings (results to be published in 2012).

Graph algorithms on complex biological networks seems to become fairly important for this group of biologists, and their problems are very nice to study but perhaps too difficult (for example the theory of graph based phylogeny seems to be much more involved that of tree based and we only hardly found very partial results).

2.4.4 Collaborations

During the period we had an active collaboration including joint papers with the following groups.

- Pr. D. Corneil, University of Toronto (Canada) and his group.
- J. Stacho, actually post-doc at Warwick University (U.K.).
- P. Crescenzi, University of Florence (Italy) and his PhD students
- R. Grossi, University of Pisa (Italy) and his PhD students
- F. Dragan, University of Kent (USA).

- We also had strong relationships with V. Chepoi (Marseille), M. Latapy (University Paris 6) and C. Paul (Montpellier).
- Eric Bapteste et Philippe Lopez biologists from Systématique, adaptation, évolution (SAE) - CNRS UMR 7138 & University Paris 6. from (University Paris 6).

2.4.5 External support

Starting in 2012, the ANR project Algopol will grant us 2 years of Post-doc support for studying complex networks in Biology.

2.4.6 Self assessment

This Gromov's notion of δ-hyperbolicity applied to graphs was very appealing, but we were not so successful to capture with it the metrics aspects of the Internet network, moreover M. Soto's results in his PHD are showing that the practical applications of δ-hyperbolicity may be limited.

The success we had for exact diameter computations on huge graphs and networks can be extended to exact computations of centers and it is worth to study how these ideas can lead to distributed algorithms. But we still have to deeply understand why these algorithms are so efficient experimentally, and in particular to understand the efficiency of the *4-sweep heuristic*.

Similarly the area of graph searching still remains a privileged topic for future research with many applications both theoretical and practical.

Our main question is : *What can be learned of the structure of a given graph using a series of successive graph traversal?*

2.5 Objective (2): Large Networks

We are aiming at studying fundamental aspects of large interaction networks enabling massive distributed storage, efficient decentralized information retrieval, quick inter-user exchanges, and/or rapid information dissemination. The study is mostly oriented towards the design and analysis of algorithms for these networks, by taking into account properties inherent to the underlying infrastructures upon which they are built. The infrastructures and/or overlays considered in the project are selected from different contexts, including communication networks (from Internet to sensor networks), and societal networks (from the Web to P2P networks).

Our main objectives are to design algorithms and to analyze algorithmic phenomena in the broad context of interaction networks. The algorithms will be based on qualitative properties, considered from the nature and/or functionality of the networks, such as bounded growth rate, low doubling dimension, minor excluding, or hyperbolic metrics.

2.5.1 Personnel

Pierre Fraigniaud, George Giakkoupis, Amos Korman, Emmanuelle Lebhar, Laurent Viennot.

2.5.2 Project-team positioning

Jon Kleinberg, from Cornell University, who proposed most of the early results about small-worlds and navigability from which GANG work started.

Dalia Malkhi and her team from Microsoft Research, Silicon Valley, have pursued studies closed to ours, but with a specific focus on storage-related issues.

David Peleg, from Weizmann Intitute, Israel, also works in this field (we have otherwise some collaborations with him).

Within INRIA, Cyril Gavoille and his team at Inria EPI Cepage has worked on subject very close to ours (that justifies our strong collaborations with Bordeaux).

2.5.3 Scientific achievements

Tractability of Internet An important objective for Gang was to test large networks against theoretical parameters known to measure tractability. We have first studied the doubling dimension of Internet. A metric has small doubling dimension if any ball of radius $2r$ can be covered by a small number of balls of radius r. A large amount of algorithms has been designed for the Internet under the assumption that the distance defined by the round-trip delay (RTT) is a metric with low doubling dimension (see e.g. Slivkins [PODC'05]). In [85], we show that Internet delays can be modeled by an inframetric, i.e. a function satisfying a relaxed version of the triangle inequality. Interestingly, the theory of doubling and bounded growth metrics can be extended to this setting. Additionally, we show that Internet delays satisfy a doubling like property that allows to design an efficient compact routing scheme.

In [65], we study the measurement of the Internet according to two graph parameters: treewidth and hyperbolicity. Both tell how far from a tree a graph is. They are computed from snapshots of the Internet released by CAIDA, DIMES, AQUALAB, UCLA, Rocketfuel and Strasbourg University, at the AS or at the router level. On the one hand, the treewidth of the Internet appears to be quite large and being far from a tree with that respect, reflecting some high degree of connectivity. This proves the existence of a well linked core in the Internet. On the other hand, the hyperbolicity (as a graph parameter) appears to be very low, reflecting a tree-like structure with respect to distances enabling additive stretch compact routing [13]. Additionally, we compute the treewidth and hyperbolicity obtained for classical Internet models and compare with the snapshots.

Spanners With the growth of Internet, lookup time in routing tables may become a bottleneck. Fortunately, there exists a theoretical solution to this: compact routing which consists in trading table size versus route efficiency (Thorup and Zwick [SPAA'01]). However, such constructions are centralized and are still far from a practical protocol. As an intermediate goal, we have considered the problem of computing sparse spanners with low stretch. All known compact routing solutions implicitly solve this simpler problem. A spanner is a subgraph preserving distances (Peleg and Ullman [PODC'87]). There exists a trade-off between the size of the spanner, i.e. its number of edges, and its stretch, i.e. how much distances are stretched in the spanner compared to the graph.

In [71], we propose a distributed algorithm for computing a sparse spanner with low stretch. Interestingly, this is the first distributed algorithm that do not require a priori knowledge of the number of nodes n in the network. This can be considered as a first towards distributed construction of compact routing scheme. In [72], we study the distributed construction of spanners with nearly additive stretch (compared to multiplicative stretch).

Motivated by the fact that neighbors are generally known in practical routing algorithms, we have introduced the notion of remote-spanner [98]. This concept is the fundamental object behind the problem of optimizing the number of links advertised in a link state routing protocol. The size of optimal remote spanners where distances are preserved is analyzed in [97].

Motivated by multipath routing, we consider in [92, 93] the construction of multipath spanner, i.e. a subgraph preserving the existence of multiple paths between any pair of vertices.

Distributed compact labeling scheme. We have solved the *ancestry problem* which was introduced more than twenty years ago by Kannan et al. [STOC '88] and is among the most well-studied problems in the field of informative labeling schemes.

Specifically, we have constructed an *ancestry labeling scheme* for n-node trees with label size $\log_2 n + O(\log\log n)$ bits, thus matching the $\log_2 n + \Omega(\log\log n)$ bits lower bound given by Alstrup et al. [SODA '03]. Motivated by applications to XML trees, which typically have extremely small depth, we also provided an ancestry labeling scheme that labels n-node trees of depth d with labels of size $\log_2 n + 2\log_2 d + O(1)$ bits. As a corollary, we also obtained an adjacency scheme that labels n-node trees of depth d with labels of size $\log_2 n + 3\log_2 d + O(1)$ bits. All our schemes assign the labels in linear time, and guarantee that any query can be answered in constant time.

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

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

Social networks. We have studied two important aspects of the acquaintance network between individuals.

The "small world phenomenon" in arbitrary networks. Revisiting the "small-world" experiments of the 60's, Kleinberg observed that individuals are very effective at constructing short chains of acquaintances between any two people, and he proposed a mathematical model of this phenomenon. In this model, individuals are the nodes of a *base graph*, the square grid, capturing the *underlying structure* of the social network; and this base graph is augmented with additional edges from each node to a few *long-range contacts* of this node, chosen according to some natural distance-based distribution. In this augmented graph, a greedy search algorithm takes only a polylogarithmic number of steps in the graph size. Following this work, several papers investigated the correlations between underlying structure and long-range connections that yield efficient decentralized search, generalizing Kleinberg's results to broad classes of underlying structures, such as metrics of bounded doubling dimension, and minor-excluding graphs.

We focus on the case of *arbitrary* base graphs. We show that for a simple long-range contact distribution consistent with empirical observations on social networks, a slight variation of greedy search, where the next hop is to a distant node only if it yields sufficient progress towards the target, requires $n^{o(1)}$ steps, where n is the number of nodes. Precisely, the expected number of steps for any source–target pair is at most $2^{(\log n)^{\frac{1}{2}+o(1)}}$. This bound almost matches the best known lower bound of $\Omega(2^{\sqrt{\log n}})$ steps, which applies to a general class of search algorithms. In the context of social networks, our result could be interpreted as: individuals may well be able to construct short chains between people regardless of the underlying structure of the social network.

Emergence of the "small world phenomenon". We have proposed a dynamical process for network evolution, aiming at explaining the emergence of the small world phenomenon, i.e., the statistical observation that any pair of individuals are linked by a short chain of acquaintances computable by a simple decentralized routing algorithm, known as greedy routing. Previously proposed dynamical processes enabled to demonstrate experimentally (by simulations) that the small world phenomenon can emerge from local dynamics. However, the analysis of greedy routing using the probability distributions arising from these dynamics is quite complex because of mutual dependencies. In contrast, our process enables complete formal analysis. It is based on the combination of two simple processes: a random walk process, and an harmonic forgetting process. Both processes reflect natural behaviors of the individuals, viewed as nodes in the network of inter-individual acquaintances.

We prove that, in k-dimensional lattices, the combination of these two processes generates long-range links mutually independently distributed as a k-harmonic distribution. We analyze the performances of greedy routing at the stationary regime of our process, and prove that the expected number of steps for routing from any source to any target in any multidimensional lattice is a polylogarithmic function of the distance between the two nodes in the lattice.

Up to our knowledge, these results are the first formal proof that navigability in small worlds can emerge from a dynamical process for network evolution. Our dynamical process can find practical applications to the design of spatial gossip and resource location protocols.

2.5.4 Collaborations

Our main collaborators on that subject during the evaluation period were David Peleg, from Weizmann Intitute, Israel, and Cyril Gavoille, from Inria EPI Cepage (Bordeaux).

2.5.5 External support

ANR Project BLANC "ALADDIN" (2007-2011); ANR Project VERSO "PROSE" (2009-2012); STREP FP7 "EULER" (2010-2013).

2.5.6 Self assessment

Most of the work described in this objective can be considered complete, in particular the *tractability of Internet*, *labelling scheme*, and *small-world related* items. The work to pursue, including the work on spanners, will be conducted in the *distributed computing* objective of the next four years.

2.6 Objective (3): Decentralized Content Distribution

2.6.1 Personnel

Yacine Boufkhad, Pierre Fraigniaud, Hoang-Anh Phan, Fabien Mathieu, Fabien de Montgolfier, Diego Perino, Laurent Viennot.

2.6.2 Project-team positioning

- Within INRIA: RAP has conducted some studies related to peer-to-peer performance. They proposed a queing system approach, which is complementary from GANG tools (graph, bandwidth-oriented, unstructured approach). CEPAGE has proposed some solution to a perfect live streaming allocation, that uses all available bandwidth, while Gang approach is mostly based on near-optimality, with simpler allocation schemes that are asymptotically optimal.
- Our work on video-on-demand systems is based on a proposal by Technicolor, which pursues some research on that field in parallel to us.
- International: following Diego Perino's participation in the European project NAPA-WiNe, other partners have also worked on issues related to unstructured live streaming, including for instance Politechnico di Torino and University of Trento.

2.6.3 Scientific achievements

P2P and de Brujin graphs During the previous evaluation period, GANG/Gyroweb has contributed to the application of de Brujin graphs to P2P systems. De Brujin graphs allow to build a structured overlay with a good trade-off between routing capabilities and average node degree. However, they have some disadvantages: they are not designed to properly handle load balancing, and performing multicast diffusion through them is not straightforward.

The contribution of Gang with respect to de Brujin graphs and P2P, mainly conducted by Pierre Fraigniaud and Hoang Anh Phan, is: an enhancement of the *de Brujin*-based DHT (D2B) in order to enhance robustness against churn and failures; the proposal of several load balancing methods that can be used to ensure fairness in degree distribution [88]; *tree-farm*, a de Brujin-based multicast algorithm [89]. Most of these contributions are gathered within the PhD Thesis of Hoang Anh Phan [6].

Bandwidth-centered analysis of distributed systems In most decentralized content delivery systems, especially P2P ones, access bandwidth is considered as the main bottleneck. Most wired high speed Internet users have a asymmetric Internet access, with a typical upload bandwidth less than 1 `Mbps`. Because of that, it is generally assumed that P2P systems cannot provide the same quality of service as more centralized systems.

In *Playing with the Bandwidth Conservation Law* [52], Diego Perino participated to a study devoted to the understanding of bandwidth dimensionning for P2P systems. One of the main highlights of the study was that using the P2P mechanism of seeders (peers that possess the content and help to its dissemination), the available bandwidth per user could be adjusted to meet almost any requirements.

This initial study was recently expanded in two directions. In [118, 103], the principle of *Live Seeding* was presented, which allows to used the seeder mechanism even when the peers that are used to *seed* the content do not initially possess it. This idea was also inspired by a patent (N. 0851025) filled by Fabien Mathieu and Diego Perino. Another, more ambitious, study is presented in [113], and extends the dimensioning framework to Future Networks where the bandwidth bottleneck is not located at the access anymore.

Live Streaming systems Live streaming is a way of distributing content that offers specific challenges: no pro-active dissemination can be used, and a low play-out delay (latency induced by the delivery mechanism) is expected. In 2008, live streaming research was roughly divided into two main categories: structured solutions, for which clean, theoretically proven, solutions were proposed, with no large scale validation, and empirical, mostly closed source, unstructured solutions, used by a large number of people.

The starting point of GANG's work on Live Streaming was a prototype developed by Diego Perino during his Internship: PULSE (*P*eer-to-Peer *U*nstructured *L*ive *S*treaming *E*xperiment). An early study conducted in 2007 allowed to start a fruitful collaboration with Thomas Bonald (France Télécom), Laurent Massoulié and Andy Twigg (Thomson/Technicolor), leading to a seminal work on epidemic Live streaming [53]. This work has then been extended: to take into account a larger class of resource balancing policies [119]; with a more subtle modeling of the chunk dissemination process, allowing more precise tuning guidelines [117]. In [11], we present and discuss possible architectures for Live P2P systems to manage overlays that try to cope with the underlying network.

Video-on-Demand Systems The challenge of distributed Video-on-Demand is to be able to manage a large catalog of videos on many, possibly heterogeneous, devices, and to be able to sustain the pattern of requests issued by users. For a request to be served, the video must obviously be available on some devices, but it is also required that those devices still possess enough available bandwidth to satisfy the QoS requirements.

Gang contribution on this subject has been inspired by the Push-to-Peer approach proposed by Thomson, where the content is pro-actively dispatched during off-peak hours. On the theoretical side, we investigated the Push-to-Peer approach, proved new results on the feasibility of a *Push-to-Peer*-like system and gave near-optimal allocation strategies [54, 57]. On a more practical side, we first developed a simulator to validate the theoretical results [56], and we started to develop a real prototype. The current version of the prototype can work without the off-peak allocation.

2.6.4 Collaborations

- Within INRIA: a collaboration is currently in progress with the EPI CEPAGE within the ANR project SONGS. The goal is to use CEPAGE expertise in simulation to perform large scale testing of some of the algorithms developed within GANG.
- Industrial: some of GANG's achievements, especially related to Live Streaming studies, come from fruitful collaborations with Orange Labs, through the CRC MARDI, and with Technicolor (informal collaboration).
- International: a strong collaboration has also been made with VTT Technical Research Center of Finland through the EIT ICT Labs and the LINCS.

2.6.5 External support

ANR SONGS: CRC MARDI; IET ICT Labs; LINCS.

2.6.6 Self assessment

The analysis of de Brujin graphs is now finished for our research team. The bandwidth-centered study of P2P systems still offers some promising challenges that should be investigated in the next four years: fine modeling of future networks, accounting for data availability issues, large scale validation. The study of Live Streaming systems is mostly finished, although a few issues should still be addressed during the year 2012. The study of Video-on-Demand should be continued, as we are currently moving from a *push-to-peer* paradigm to a *lazy allocation* paradigm that stills need some work on the theoretical part.

For all the issues that remain to be investigated within this objective, we aim at providing large scale validation through simulation, emulation and prototypes. We also propose a shift in the field of applications: in 2008, P2P systems were almost the only proposal for achieving large scale, decentralized, content distribution. The current trend is in the multiplication of proposals for decentralized content distribution, some of them already deployed (Content Distribution Networks, Cloud), others on long term (Content-Centric Networking). We think that our current expertise also apply to these systems, for which we plan to propose solutions based on P2P results.

2.7 Objective (4): Distributed computing

Large-scale networks typically include several thousands basic computing elements endowed with communication capabilities. Because of the large number of involved components configuration changes, failures (permanent, transient, or intermittent) and malicious attacks in such systems are frequent.

However, in the systems of such scale, the combined system information is so large that the resources of a single node are insignificant compared to the resources required to manage the entire system (e.g. store and process it). Also due to the scale, human intervention for fault correction and security defense is either impossible or ineffectual. At large scales, the effective tolerance to faults changes, and malicious behavior is not well understood.

Our objective is to study model of failures for such systems and the conditions under which it is possible to implement robust services.

Traditional fault detection mechanisms usually ignore mobility and autonomy constraints found in sensor networks, and the extreme volatility of peer to peer networks. A fault model should (a) accurately represent the physical phenomena occurring in practical systems (b) allow simple reasoning and efficient handling of faults. A variety of detrimental system behavior (gradual degradation of system resources for example) can be modeled as faults. We propose to focus on selecting appropriate fault models to represent complex system faults.

The consensus problem is recognized as fundamental in the design and implementation of reliable asynchronous distributed systems in presence of process failures: most of agreement problems can be reduced to consensus. Unfortunately, a classical result from Fisher, Lynch and Paterson proves that consensus has no deterministic solution in asynchronous systems even if only a single process is allowed to crash. However, the asynchronous system model is quite convenient for system designers as it does not consider a bound on the difference in processing speed of processes and message propagation. Thus this model is applicable in arbitrary computing environment. A number of ways to circumvent this impossibility have been proposed: probabilistic solution, partially synchronous system, failure detectors....

2.7.1 Personnel

Julien Clement, Carole Delporte, Hugues Fauconnier, Pierre Fraigniaud, Hung Tran-The, Laurent Viennot.

2.7.2 Project-team positioning

We focus on the fault tolerance aspects of distributed computing. We mainly deal with robust application. Fault tolerance is also studied by Inria Regal with the point of view of transient fault (self stabilization).

We take advantage of expertise of the D-NET Team and ASAP team in large scale networks in order to be familiar to the fault tolerance problem in such networks.

For the work on failure detectors we are in competition with several teams (Michel Raynal (ASAP), Petr Kuznetsov (Deutsche Telecom), Sergio Rajsbaum(Univ Mexique), Eli Gafni (UCLA) ... Our collaboration with Rachid Guerraoui and Sam Toueg (one of the conceptors of failure detectors) gives us a very strong position among these teams.

2.7.3 Scientific achievements

Failure detectors A number of ways to circumvent the impossibility to achieve consensus in an asynchronous system with failure have been proposed. Chandra and Toueg [CT96a] introduced the concept of unreliable failure detectors. A failure detector is an oracle that provides each process with information about network crashes. In general this information is the list of processes that are suspected to have crashed. The failure detector thus allows to encapsulate the synchrony necessary to solve consensus in an otherwise purely asynchronous system. A failure detector can make mistakes by not suspecting a crashed process or by erroneously suspected a correct process. Various failure detector classes can be defined and implemented depending on the properties of information they provide. One of the attractive features of the concept of failure detector is the fact that it is possible to define the weakest failure detector enabling to solve a given problem (e.g. consensus). In [17], we determine the weakest failure detectors to implement shared atomic objects in a distributed system with crash-prone processes. We first determine the weakest failure detector for the basic register object. And then the weakest failure detector for consensus in every environment. This answers to a question that has been open since 15 years. We then use that to determine the weakest failure detector for all popular atomic objects including test-and-set, fetch-and-add, queue, consensus and compare-and-swap, which we show is the same.

New model of faults In systems such as Pastry[1] or Chord[2], assuming that all processes have unique (unforgeable) identifiers might be too strong an assumption in practice. We may wish to design protocols that still work if, by a rare coincidence, two processes are assigned the same identifier.

So we have considered a distributed system with Byzantine failures in which ℓ distinct identifiers are assigned to n processes, where $1 \leq \ell \leq n$. Several processes may be assigned the same identifier, in which case we call the processes *homonyms*. This model generalizes the classical scheme where processes have distinct identifiers (i.e., $\ell = n$), and the less classical scheme where processes are anonymous (i.e., $\ell = 1$).

This approach is also useful if security is breached and a malicious process can forge the identifier of a correct process, for example by obtaining the correct process's private key. This approach also apply when users of a system may wish to preserve their privacy by remaining anonymous. In this case several processes share the same identifiers.

Unfortunately, in a fully anonymous system where no identifiers are used, very few problems are solvable. We have shown necessary and sufficient conditions in which byzantine agreement is solvable. These results have been published in [70]. Some questions remain open on this topic, for example what happen if the byzantine process may also forge the identifier of some correct process.

Verification Population protocols of Angluin *et al.* are used as a theoretical model for a collection (or population) of tiny mobile agents that interact with one another to carry out a computation. We verify our algorithms by model-checking the basic population protocol model of Angluin *et al.* This problem has received special attention during the last two years and new tools have been proposed to deal with it. We show in [60, 61] that the problem can be solved using the existing model-checking tools, e.g., Spin and Prism. For this, we apply the counter abstraction to obtain an abstraction of the population protocol model which can be efficiently verified by the existing model-checking tools. Moreover, this abstraction preserves the correct stabilization property of population protocol models. To deal with the fairness assumed by the population protocol models, we provide two new

[1]A. Rowstron and P. Druschel (Nov 2001). "Pastry: Scalable, decentralized object location and routing for large-scale peer-to-peer systems". *IFIP/ACM International Conference on Distributed Systems Platforms (Middleware), Heidelberg, Germany*: 329–350.

[2]Stoica, Ion et al. (2001). "Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications". *Proceedings of SIGCOMM'01 (ACM Press New York, NY, USA).*

recipes. The first one gives sufficient conditions under which the population protocol model fairness can be replaced by the weak fairness implemented in Spin. We show that this recipe can be applied to several population protocol models. In the second recipe, we show how to use the probabilistic model-checking and the tool Prism to deal completely with the fairness of the population protocol models. The correctness of this recipe is based on existing theorems on finite discrete Markov chains.

2.7.4 Collaborations

In LIAFA, Modeling and Verification team address the development of algorithmic approaches to system verification, from theoretical foundations to innovative verification tools. We have worked with Mihaela Sighireanu (LIAFA) from this team, who is a specialist of model checking. The result of this collaboration is detailed in the achievement Verification ([60, 61]).

For many years, we have collaborated with Stéphane Devismes (VERIMAG), who has a strong expertise in auto-stabilizing approach of fault tolerance (publications with him: [15, 66, 67]).

We collaborate closely with two Inria projects : Inria Regal, Inria ASAP (and in particular with Anne-Marie Kermarrec) in the ANR Verso Shaman. We will continue our collaboration with Inria ASAP in the ANR Displexity that begins in january 2012.

We have a strong collaboration with Rachid Guerraoui (EPFL Lausanne) and Sam Toueg (Univ Toronto Canada) for several years. From 2008, we have published 4 papers in top level conferences and 2 journal papers (Distributed computing and J. ACM) with Rachid Guerrraoui and one paper in top level conference paper and 2 journal papers (Distributed computing) with Sam Toueg. We have also published others papers with them.

We collaborate also Marcos Aguilera(from Microsoft research, Palo Alto, USA), Mikel Larrea (from University of the Basque Country) and Eric Ruppert (from York university, Canada).

With Anne-Marie Kermarrec, Rachid Guerraoui and Eric Ruppert we worked on the achievement “New model of faults”, we have had one publication in a top level conference and we are invited to publish our paper in distributed computing([69, 70]).

With Sam Toueg and Marcos Aguilera, we work on partially synchronous model, we have had one publication in a top level conference and a journal paper (Distributed computing DFG11).

With Stephane Devismes and Mikel Larrea, we work also in partially synchronous model ([67]).

2.7.5 External support

- PICS(2010-2013) CNRS International Projects for Scientific Cooperation - with LPD (EPFL)
- ANR Verso Shaman: the two others teams are Inria Regal/grand Large, Inria Asap
- ANR Blanche Displexity: led by GANG, the two others teams are LABRI, Inria Asap
- PEFICAMO -Ile de France project. Coordinator: LIAFA (P7) - Partners : LSV (ENS-C) - LPD (EPFL)

2.7.6 Self assessment

We have obtained strong result concerning the failure detectors approach. We think that most of the works on weakest failure detectors has been done, although some questions are still open. These questions seem very hard to solve (for example the weakest failure detector in message passing for k-set agreement). In the next four years, we want to focus more on the failure model *for large scale systems* and, more broadly, to address issues related to the foundations of distributed computing.

3 Knowledge dissemination

3.1 Publications

	2008	2009	2010	2011
PhD Thesis	1	1	2	2
H.D.R (*)				
Journal	3	10	13	14
Conference proceedings (**)	6	19	15	20
Book chapter	1			
Book (written)				
Book (edited)			1	1
Patent	1			
General audience papers		3		
Technical report		3	2	3
Deliverable				

(*) HDR Habilitation à diriger des Recherches
(**) Conference with a program committee

Indicate the major journals in the field and, for each, indicate the number of papers coauthored by members of the project-team that have been accepted during the evaluation period.

1. Algorithmica: 4 ([14, 24, 48, 13])
2. Discrete Applied Mathematics: 2 ([12, 20])
3. Distributed Computing: 6 ([37, 47, 9, 18, 19, 23])
4. Journal of the ACM: 1 ([17])
5. Pure and Applied Mathematics: 3 ([30, 26, 27])
6. Theoretical Computer Science: 6 ([36, 46, 38, 44, 42, 21])

Indicate the major conferences in the field and, for each, indicate the number of papers coauthored by members of the project-team that have been accepted during the evaluation period.

1. DISC: 4 ([68, 72, 75, 90])
2. FOCS: 1 ([84])
3. ICALP: 2 ([58, 74])
4. ICC: 2 ([89, 108])
5. ICDCS: 1 ([60])
6. INFOCOM: 1 ([85])
7. IPDPS: 2 ([57, 98])
8. P2P: 2 ([52, 103])
9. PODC: 7 ([50, 70, 71, 76, 78, 101, 102])

10. SAT: 2 ([49, 55])
11. Sigmetrics: 1 ([53])
12. SIROCCO: 3 ([51, 67, 92])
13. SODA: 2 ([77, 82])
14. SPAA: 4 ([69, 79, 86, 81])
15. STACS: 1 ([99])
16. STOC: 3 ([80, 83, 109])
17. SSS: 1 ([66])

3.2 Software

PULSE

PULSE was a P2P live streaming client developed by Diego Perino following a design proposed by Fabio Pianese. Its code is available on

`http://napa-wine.eu/cgi-bin/twiki/view/Public/PULSE`

PULSE is not maintained anymore. However, its code served as a starting point for another client developped within the NAPA-WiNe project: PeerStreamer

`http://napa-wine.eu/cgi-bin/twiki/view/Public/PeerStreamer`

VoD Prototype

The prototype was developed to test ideas born from [54, 57, 56] in realistic conditions. It currently contains many improvements compared to the original design, especially on the cache management side (original design relied on a pro-active cache filling that is not necessary anymore). Although not being made public yet, mainly because of a lack of resources, the prototype is promising, and it is expected to grow in maturity thorough the ANR INFRA SONGS project.

Palabre/Peerple : a Cooperative Peer-to-peer Platform

The peer-to-peer paradigm can be used to duplicate sensible data. GANG developped a peer-to-peer client for personal files sharing and backup. The Peerple project is not active anymore, but led to the creation of the Move 'N Play company . Peerple is an open-source project http://peerple.gforge.inria.fr/.

Simulators

Epidemic Live Streaming The work on epidemic live streaming that started with [53] led to the development of chunk exchange simulator that was used in subsequent articles [96, 117]. This tool is based on the research conducted with Orange Labs and Technicolor Research, and also served as a guideline for the development by Politecnico di Torino of a more versatile simulator, P2PTVSim.

P2P Future Networks The on-going work on P2P filesharing in future networks [113] considers peers of a P2P system as a point process. A custom-made simulator has been designed to support the study.

3.3 Valorization and technology transfert

Anh-Tuan Gai, formerly a postdoctoral fellow in the Peerple project, has created Move 'N Play company for disseminating this software (scope: distant access to personal media content). Move 'N Play won an award of best innovative entreprise.

3.4 Teaching

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

Carole Delporte is: in charge of a course entitled "Distributed programming" (Professional Master, Paris Diderot University, 192 hours); in charge of "Algorithmique distribuée avec mémoire partagée" (MPRI, University of Paris Diderot).

Hugues Fauconnier is: in charge of both courses "Internet Protocols" and "Distributed algorithms" (Professional Master, Paris Diderot University, 192 hours); in charge of "Algorithmique distribuée avec mémoire partagée" (MPRI, University of Paris Diderot).

Pierre Fraigniaud is in charge of a course entitled "Algorithmique distribuée pour les réseaux" (MPRI, University of Paris Diderot, 12 hours).

Michel Habib is in charge of: a course entitled "Graph Algorithms" (MPRI, University of Paris Diderot); also "Graph Algorithms" in L3 (Ens Cachan), "Advanced Discrete Algorithms" M1 Paris Diderot and also two courses untitled: Search Engines; Parallelism and mobility, which includes peer-to-peer overlay networks (Professional Master, Paris Diderot University, 192 hours).

Fabien Mathieu is teaching Peer-to-peer Networks (Master 2 Computer Science, University of Paris 6, 6 hours; ENS Lyon, L3, 2 hours).

Fabien de Montgolfier is teaching: foundation of computer science, algorithmics, and computer architecture (Computer Science U.F.R., University of Paris Diderot, 192 hours); P2P theory and application (University of Marne-la-Vallée);

Laurent Viennot is teaching at MPRI, University of Paris Diderot: ad hoc and web graph algorithms in the "networks dynamics" course (2008–2009); "Networks and geometry" (2010); "Structures de données distribuées et routage" (2011, 12 hours).

3.4.1 Ongoing theses

- Heger Arfaoui *Distributed Computational Complexity*
- Hervé Baumann *protocoles d'échange dans les réseaux distribués.*
- Xavier Koegler *Population protocols*
- Antoine Mamcarz *Algorithmes de décomposition pour les grands graphes* (CNRS grant)
- Hung Tran: *Failure detection with Byzantine adversary* (ANR VERSO Shaman)
- Jérémie Dusart: *Graph searches* (BMN)

3.4.2 Defended theses

- Thomas Hugel *Estimations de satisfaisabilité.* Defended December 7th 2010.
- Vincent Limouzy *Algorithmes de décomposition de graphes* (MENRT). Defended December 3rd, 2008.
- Diego Perino *On Resource allocation algorithms for peer-to-peer multimedia streaming.* (CIFRE Orange Labs). Defended November 16th 2009.
- Hoang Anh Phan *Equilibrage de charge et diffusion multicast dans les systèmes pair-à-pair.* Defended September 29th 2010.
- Mauricio Soto *Quelques propriétés topologiques des graphes et applications à Internet et aux réseaux.* Defended on December 2nd 2011.
- Thu-Hien To *On some graph problems in phylogenetics.* Defended on September 15th 2011.

3.5 General Audience Actions

- Fabien de Montgolfier and Vincent Limouzy published in)i(nterstices a popularization of the tree-decomposition of graphs[3]
- Laurent Viennot published in)i(nterstices/docsciences popularization of networks[4]
- M. Habib gave a few invited lectures:
  - *Efficient graph algorithms*, VI Escuela de Verano en Matemáticas Discretas, Valparaiso, Chili, janvier 2011;
  - at Jiao Tong University Shanghai (University of the famous Ranking), China : 2 lectures on modular decomposition and chordal graphs, october 2011.

3.6 Visibility

Carole Delporte is Program Committee member of: OPODIS 2010; ICDCN 2012; ICDCS 2012; PODC 2012; Algotel 2011.

Dominique Fortin is a regular reviewer for international journals in optimization: Applied Math. Modeling, Computers & Operations Research, European Journal of Operational Research, Mathematical Reviews (since 2011).

Hugues Fauconnier is Program Committee member of: PODC 2010; SSS 2010; ICDCS 2012; Algotel 2010; Algotel 2012.

Pierre Fraigniaud is:

- Technical Program Chair (TPC) of PODC 2011 (with Cyril Gavoille).
- Keynote invited speaker at: PODC 2010; ICALP 2010; ICDT 2010.
- Program Committee member of: PODC 2009; MFCS 2009; ICALP 2009; ESA 2009; IPDPS 2009; ICDCN 2009; SSS 2009; DIALM-POMC 2008; ICDCS 2008.

[3] `http://interstices.info/pacman`
[4] `http://interstices.info/reseaux-fibre-information`
`http://interstices.info/idee-recue-informatique-10`

- Steering Committee member of: STACS; SPAA; DISC.
- Editorial Board member of: Distributed Computing; Theory of Computing Systems; Journal of Interconnection Networks.

Michel Habib is:

- member of the steering committee of STACS (Symposium on Theoretical Aspects of Computer Science) and also of WG (International Workshop on Graph-Theoretic Concepts in Computer Science).
- Program Committee member of: ICALP 2009; MFCS 2010; SWAT 2012.

Fabien Mathieu is:

- member of the LINCS board committee.
- TPC of: Algotel 2012 (with Nicolas Hanusse; MARDI 2007 workshop.
- General Chair of Rescom *journées non-thématiques* 2011 (with Lélia Blin).
- PC member of: ITC 2012; IEEE IPDPS 2011; JDIR 2008 and 2010; SSS 2009.
- Regular reviewer for international journals and conferences (IEEE Transactions on Information Theory, Transactions on Computer, JSAC, Opodis, Euro-Par, ...).

Laurent Viennot is a scientific editor of the)i(nterstices (http://interstices.info/) vulgarization site initiated by Inria in collaboration with french universities and Cnrs. He has written an article on the differences between the web and internet [111].

4 External Funding

(k€)	2008	2009	2010	2011
National initiatives				
ANR GRAAL	33	33		
ANR ALADDIN	40	60	60	40
ANR Shaman	27	36	35	19
ANR Prose			10	20
European projects				
e.g. FP7 XX				
FP7 EULER			13	7
EIT ICT				15
Industrial contracts				
Orange Labs CRC MARDI	42	42		
Other funding				
PEFICAMO		10	55	
PICS				7
BQR				12
Total	142	181	173	120

National initiatives

ANR ALADDIN (Algorithm Design and Analysis for Implicitly and Incompletely Defined Interaction Networks) : Pierre Fraigniaud is leading an ANR project "blanc" (i.e. fundamental research) about the fundamental aspects of large interaction networks enabling massive distributed storage, efficient decentralized information retrieval, quick inter-user exchanges, and/or rapid information dissemination. The project is mostly oriented towards the design and analysis of algorithms for these (logical) networks, by taking into account proper ties inherent to the underlying infrastructures upon which they are built. The infrastructures and/or overlays considered in this project are selected from different contexts, including communication networks (from Internet to sensor networks), and societal networks (from the Web to P2P networks). Ending in november 2011, the project is prolonged until end of 2012 for LABRI partner. Aladdin granted 1 postdoc for GANG (George Giakkoupis).

ANR Displexity Managed by University Paris Diderot, C. Delporte and H. Fauconnier lead this project that grants 1 Ph. D. and 2 internships per year. The main goal of DISPLEXITY(*calcul DIStribué : calculabilité et comPLEXITé*) is to establish the scientific fundations of a theory of calculability and complexity for distributed computing. Partners of GANG are IRISA and LABRI. Displexity starts in 2012.

ANR Graph Decomposition and Algorithms (GRAAL) is an ANR project "blanc" (i.e. fundamental research) about graph decompositions with LABRI and LIRMM. This project deals with fundamental aspects of computer science, namely theoretical and algorithmic aspects of decomposition methods for graphs and various combinatorial structures and extensions such as matroids, countable graphs... It proposes to combine approaches issued from graphs, discrete algorithms, formal languages and logic.

ANR INFRA-SONGS aims at designing a unified and open simulation framework for next generation systems. This framework should allow to study the four following domains: Grids, Peer-to-Peer systems, Clouds and High Performance Computation systems. Fabien Mathieu participates to this project in collaboration with the CEPAGE EPI. Project starts in 2012.

ANR PROSE 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. Pierre Fraigniaud is the local coordinator of this project.

ANR 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. Hugues Fauconnier is local manager for this project, which granted Ph. D. H. Tran.

LINCS GANG participates in the Laboratory of Information, Networking and Communication Sciences (LINCS); http://www.lincs.fr/ created on October 28th, 2010, by three French institutions of higher education and research and officially launched in 2011. LINCS institutional members are: INRIA, Institut Télécom and UPMC. Alcatel-Lucent joined the LINCS in February 2011 as a strategic partner.

European projects

Experimental UpdateLess Evolutive Routing (EULER) is a 3-year STREP Project targeting Challenge 1 "Technologies and systems architectures for the Future Internet" of the European Commission (EC) Seventh Framework Programme (FP7). The project scope and methodology position within the FIRE (Future Internet Research and Experimentation) Objective ICT-2009.1.6 Part b: Future Internet experimentally-driven research . The project has 7 partners: INRIA, Alcatel, UPMC, UCL, IBBT, RACTI, CAT. `http://www.euler-fire-project.eu/`.

EIT ICT Labs Project Fundamentals of Networking (FUN) (action line *Internet Technologies and Architectures*) 2011; Partners: INRIA, VTT, Aalto University, Eindhoven University. `http://eit.ictlabs.eu/ict-labs/all-events/article/fundamentals-of-future-networking-workshop/`. 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. The project allowed to invite Pascal Felber to work with GANG on P2P modeling for one month.

COST 295 – Dynamo is an action of the European COST program (European Cooperation in the Field of Scientific and Technical Research) inside of the Telecommunications, Information Science and Technology domain (TIST). It is leaded by Pierre Fraigniaud (Chair of the managing committee). It gather more than 30 sites all over Europe around Dynamic Communication Networks. The Action is motivated by the need to supply a convincing theoretical framework for the analysis and control of all modern large networks. This will be induced by the interactions between decentralised and evolving computing entities, characterised by their inherently dynamic nature.

Industrial contracts

CRC MARDI (2006–2010) was a collaboration contract between Inria (GANG) and Orange Labs France Telecom. The subject was the study of decentralized networks over Internet. The grant resulted in the PhD funding (CIFRE contract) of Diego Perino, co-supervised by Fabien Mathieu and Laurent Viennot.

ADR LINCS is an Alcatel grant in which GANG gets 25 kE /y, starting in 2012. The funding will be used to work on the subject *manipulability of voting systems and applications to networks.*

A CIFRE (PhD funding) ALCATEL is starting in 2012. A PhD student (The Dang Huynh) will work on *Extension of PageRank for Social Networks.*

Other funding,

PEFICAMO (Performance and Reliability of Mobile Sensor Networks) is a project from *Région Île de France* managed by Hugues Fauconnier, granting Julien Clément. It aims at understanding the behavior of sensor networks, and designing efficient algorithms adapted to these networks. Other partners are ENS Cachan and EPFL.

Architectures multicoeurs et algorithmique distribuée (2011–2013) is a PICS grant (CNRS funding for international cooperation) between Rachid Guerraoui (EPFL-Suisse) and members of GANG. 7k€ were granted in 2011.

Algorithmique distribuée quantique is a BQR (*Bonus Qualité Recherche*, university funding) that granted 12k€ in 2011.

5 Objectives for the next four years

Taking into account the scientific achievements of the last four years, GANG has decided to focus for the next four years on the following objectives:

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

5.1 graph algorithms

Graph Decompositions We will 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 somme well-known theorems such as the Perfect Graph Theorem. For these decompositions there is a lack of efficient algorithms. We will try to design algorithms working in $O(nm)$ since we think that this could be a lower bound for these decompositions.

Graph Search We will focus our research on graph search to *multi-sweep graph searches.* In this domain a graph search only yields a total ordering of the vertices which can be used by the following 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 the quadratic assignment problems).

Application to exact computations of centers in huge graphs, but also to maximal flow problems.

Complex networks in Biology We will continue our collaboration with E. Bapteste et P. Lopez and their group form Systématique, adaptation, évolution (SAE) - CNRS UMR 7138 & University Paris 6.

This collaboration will develop ad-hoc software tools for analyzing graphs coming form biology.

5.2 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.) (C. Delporte et H. Fauconnier) while the other one is mostly interested in spatial issues (network structure, memory requirements, etc.)(P. Fraigniaud, A. Korman, L. Viennot). 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.

We conduct this project with two other teams: in LABRI and ASAP in the frame of the ANR blanche DISPLEXITY.

The main question addressed in this project is the following: 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 [5], 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 [6] as well as Linial's lower bound [7] 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 $\mathcal{LOCAL}$ model, which, though they do not claim to reflect accurately real distributed computing systems, enable focusing on some core issues.

Despite the above issues, this project seeks to demonstrate that many important notions of Distributed Computing seem to fit well with standard computational complexity. Distributed Computing should thus greatly benefit from expressing its main challenges in this standard framework for making them accessible to a wider audience.

5.3 A peer-to-peer approach to Future Content Distribution

With the increase of content distribution through the Internet, it is expected that this field of applications will continue to present interesting research challenges in the years to come, so GANG aims at pursuing the research in this domain.

Streaming prototyping As a result of the NAPA-WiNe project, where Diego Perino was involved in, a live streaming prototype based on PULSE was designed: PeerStreamer. Its main feature is the possibility to change the diffusion algorithm at will, and in particular to use one of the algorithms that GANG has proposed in the last years.

---

[5] yes-no problems are generally called decision problems in classical computing theory. They are problems that output yes or no depending on the inputs of the nodes. A precise definition of yes-no problems in distributed computing is not obvious. Yes-no problems is the subject of Task 2.

[6] FLP result is a fundamental result in fault-tolerant distributed computing that proves the impossibility of consensus when at least one process may crash. Reference: Fischer, Michael J; Nancy A. Lynch, Michael S. Paterson (April 1985). "Impossibility of distributed consensus with one faulty process".*Journal of the ACM 32 (2).* doi:10.1145/3149.214121.

[7] Linial, Nathan (1992). "Locality in Distributed Graph Algorithms". *Siam J. Comput., vol. 21n Nr. 1.* p. 193–201.

Unrelated to that, resulting from the collaboration with Orange Labs on distributed VoD algorithms, GANG has also designed an open source VoD prototype, with a modular design that allows to test a wide range of use cases and algorithms.

Both prototypes have a good potential, and it is planned to begin new experiments, among other things to validate new algorithms. Such work is planned mostly through the ANR SONGS project (2012–2015), but also by using local testbeds available at the LINCS premises.

New network models The new models that have been proposed to take into account the evolution of network architecture and usage [113] 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 (the experimental part is also scheduled within the ANR SONGS project).

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, could lead to some significant results in a short term perspective, and a collaboration to address them has recently begun with Pascal Felber (Neuchâtel).

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, for instance at Orange Labs, Alcatel, or at the INRIA EPI RAP, but most of them rely on operational research methodology and offer solutions that can sometimes be difficult to use in practice. GANG plans to use a different approach, based on alternate modeling assumptions inherited from our previous work on bandwidth dimensioning [52]. The goal of this complementary approach is to provide simple dimensioning guidelines while giving approximated, yet meaningful, performance evaluation.

Long term perspective on P2P content distribution

Clearly, P2P is currently not as popular as it was a few years ago, both for the public and the research communities. Many reasons exist: growth of new Internet uses (YouTube, DailyMotion, ...), emergence of new research themes...

However, we think that the interest of basic research in P2P is not questioned by these changes. For instance, the success of YouTube-like delivery platforms does not (only) come from its technical performance[8], 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 that 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 can also change the power balance very fast. For example, the recent shutdown of the main semi-centralized filesharing service (MegaUpload) is causing unexpected changes in the content distribution ecosystem, with effects that are still to determine.

For all these reasons, it seems important to continue to build on GANG existing expertise in the field of decentralized content distribution, even if it is quite difficult at the present time to tell if that expertise should apply on traditional P2P, CCN or Cloud architectures, or on any hybridation of these.

[8]For instance most of Youtube videos are played using the Adobe Flash Player, which is argued not to be a very smart choice with respect to performance.

6 Bibliography of the project-team

Books and Monographs

[1] Cyril Gavoille and Pierre Fraigniaud, editors. *Proceedings of the 30th Annual ACM Symposium on Principles of Distributed Computing, PODC 2011, San Jose, CA, USA, June 6-8, 2011.* ACM, 2011.

[2] Christophe Paul and Michel Habib, editors. *Graph-Theoretic Concepts in Computer Science, 35th International Workshop, WG 2009, Montpellier, France, June 24-26, 2009. Revised Papers*, volume 5911 of *Lecture Notes in Computer Science*, 2010.

Doctoral dissertations and "Habilitation" theses

[3] Thomas Hugel. *Estimation de satisfaisabilité.* PhD thesis, Paris Diderot, Dec 2010.

[4] Vincent Limouzy. *Algorithmes de décomposition de graphes.* PhD thesis, Paris Diderot, Dec 2008.

[5] Diego Perino. *On Resource allocation algorithms for peer-to-peer multimedia streaming.* PhD thesis, Paris Diderot, Nov 2009.

[6] Hoang Anh Phan. *Equilibrage de charge et diffusion multicast dans les systèmes pair-à-pair.* PhD thesis, Paris Diderot, Dec 2010.

[7] Mauricio Soto. *Quelques propriétés topologiques des graphes et applications à Internet et aux réseaux.* PhD thesis, Paris Diderot, Dec 2011.

[8] Thu-Hien To. *Quelques problèmes de graphes en phylogénie.* PhD thesis, Paris Diderot, Sep 2011.

Articles in referred journals and book chapters

[9] Hervé Baumann, Pierluigi Crescenzi, and Pierre Fraigniaud. Parsimonious flooding in dynamic graphs. *Distributed Computing*, 24(1):31–44, 2011.

[10] Anne Bergeron, Cedric Chauve, Fabien de Montgolfier, and Mathieu Raffinot. Computing common intervals of k permutations, with applications to modular decomposition of graphs. *SIAM J. Discrete Math.*, 22(3):1022–1039, 2008.

[11] Robert Birke, Emilio Leonardi, Marco Mellia, Arpad Bakay, Tivadar Szemethy, Csaba Kiraly, Renato Lo Cigno, Fabien Mathieu, Luca Muscariello, Saverio Niccolini, Jan Seedorf, and Giuseppe Tropea. Architecture of a network-aware P2P-TV application: The napa-wine approach. *IEEE Communications Magazine*, 49(6):154–163, 2011.

[12] Yacine Boufkhad and Thomas Hugel. Estimating satisfiability. *Discrete Applied Mathematics*, 160(1-2):61 – 80, 2012.

[13] Victor Chepoi, Feodor Dragan, Bernard Estrellon, Michel Habib, Yann Vaxes, and Yang Xiang. Additive spanners and distance and routing labeling schemes for hyperbolic graphs. *Algorithmica*, on line first, 2011.

[14] Reuven Cohen, Pierre Fraigniaud, David Ilcinkas, Amos Korman, and David Peleg. Labeling schemes for tree representation. *Algorithmica*, 53(1):1–15, 2009.

[15] Carole Delporte-Gallet, Stéphane Devismes, and Hugues Fauconnier. Stabilizing leader election in partial synchronous systems with crash failures. *J. Parallel Distrib. Comput.*, 70(1):45–58, 2010.

[16] Carole Delporte-Gallet and Hugues Fauconnier. Objets partagés et détecteurs de défaillances. *Technique et Science Informatiques*, 30/7:841–871, 2011.

[17] Carole Delporte-Gallet, Hugues Fauconnier, and Rachid Guerraoui. Tight failure detection bounds on atomic object implementations. *J. ACM*, 57(4), 2010.

[18] Carole Delporte-Gallet, Hugues Fauconnier, Rachid Guerraoui, and Andreas Tielmann. The disagreement power of an adversary. *Distributed Computing*, 24(3-4):137–147, 2011.

[19] Carole Delporte-Gallet, Hugues Fauconnier, and Sam Toueg. The minimum information about failures for solving non-local tasks in message-passing systems. *Distributed Computing*, On line first, 2011.

[20] Arnaud Durand and Michel Habib. Complexity issues for the sandwich homogeneous set problem. *Discrete Applied Mathematics*, 159(7):574–580, 2011.

[21] Yuval Emek, Pierre Fraigniaud, Amos Korman, and Adi Rosén. Online computation with advice. *Theor. Comput. Sci*, 412(24):2642–2656, 2011.

[22] Yuval Emek, Pierre Fraigniaud, Amos Korman, and Adi Rosén. On the additive constant of the k-server work function algorithm. *Information Processing Letters*, 110(24):1120–1123, 2010.

[23] Yuval Emek and Amos Korman. New bounds for the controller problem. *Journal of Distributed Computing*, 24(3-4):177–186, 2011.

[24] F. V. Fomin, P. Fraigniaud, and N. Nisse. Nondeterministic graph searching: From pathwidth to treewidth. *Algorithmica*, 53(3):358–373, 2009.

[25] Dominique Fortin. *Optimal Control: Theory and Applications*, chapter Hadamard's matrices, Grothendieck's constant and square Root Two. Springer, 2009.

[26] Dominique Fortin. Eigenvectors of Toeplitz matrices under higher order three term recurrence and circulant perturbations. *Int. J. Pure Appl. Math.*, 60(2):217–228, 2010.

[27] Dominique Fortin. Eigenvectors of three term recurrence Toeplitz matrices and Riordan group. *Int. J. Pure Appl. Math.*, 67(1):97–105, 2011.

[28] Dominique Fortin and Ider Tseveendorj. Piecewise-convex maximization approach to multiknapsack. *Optimization*, 58(5-6):1–18, 2008.

[29] Dominique Fortin and Ider Tseveendorj. *Global Optimization : Theory, Methods & Applications I*, volume 12(B) of *lecture notes in decision sciences*, chapter Escaping from a local solution region of convex maximization problem, pages 943–948. Global-Link, september 2009.

[30] Dominique Fortin and Ider Tseveendorj. A trust branching path heuristic for permutation problems. *International Journal of Pure and Applied Mathematics*, 56(3):329–343, 2009.

[31] Dominique Fortin and Ider Tseveendorj. A trust branching path heuristic for zero-one programming. *European Journal of Operational Research*, 197(2):439–445, 2009.

[32] Dominique Fortin and Ider Tseveendorj. Piece adding technique for convex maximization problems. *Journal of Global Optimization*, 48:583–593, 2010. 10.1007/s10898-009-9506-z.

[33] Dominique Fortin and Ider Tseveendorj. Attractive force search algorithm for piecewise convex maximization problems. *Optimization Letters*, pages 1–17, September 2011.

[34] Dominique Fortin and Ider Tseveendorj. Piecewise convex formulations of binary and permutation problem. *Siberian Journal of Numerical Mathematics*, 14(4):409–423, 2011.

[35] Dominique Fortin and Ider Tseveendorj. Piecewise convex maximization problems: piece adding technique. *J. Optim. Theory Appl.*, 148(3):471–487, 2011.

[36] P. Fraigniaud, C. Gavoille, A. Kosowski, E. Lebhar, and Z. Lotker. Universal augmentation schemes for network navigability. *Theoretical Computer Science*, 2008.

[37] Pierre Fraigniaud, Cyril Gavoille, David Ilcinkas, and A. Pelc. Distributed computing with advice: Information sensitivity of graph coloring. *Distributed Computing*, 21:395–403, 2009.

[38] Pierre Fraigniaud, Cyril Gavoille, Adrian Kosowski, Emmanuelle Lebhar, and Zvi Lotker. Universal augmentation schemes for network navigability. *Theor. Comput. Sci*, 410(21-23):1970–1981, 2009.

[39] Pierre Fraigniaud, David Ilcinkas, and Andrzej Pelc. Communication algorithms with advice. *Journal of Computer and System Sciences*, 76(3-4), 2010.

[40] Pierre Fraigniaud, Amos Korman, and Emmanuelle Lebhar. Local mst computation with short advice. *Theory of Computing Systems*, 47(14):920–933, 2010.

[41] Pierre Fraigniaud, Emmanuelle Lebhar, and Zvi Lotker. A lower bound for network navigability. *SIAM Journal on Discrete Mathematics*, 24(1):72–81, 2010.

[42] Pierre Fraigniaud, Emmanuelle Lebhar, and Zvi Lotker. Recovering the long-range links in augmented graphs. *Theor. Comput. Sci.*, 411(14-15):1613–1625, 2010.

[43] Michel Habib and Christophe Paul. A survey of the algorithmic aspects of modular decomposition. *Computer Science Review*, 4(1):41–59, 2010.

[44] Marcos Kiwi, Mauricio Soto, and Christopher Thraves. Adversarial queuing theory with setups. *Theor. Comput. Sci.*, 410(8-10):670–687, 2009.

[45] Amos Korman. Labeling schemes for vertex connectivity. *ACM Transactions on Algorithms*, 6(2), 2010.

[46] Amos Korman and Shay Kutten. A note on models for graph representations. *Theor. Comput. Sci.*, 410(14):1401–1412, 2009.

[47] Amos Korman, Shay Kutten, and David Peleg. Proof labeling schemes. *Distributed Computing*, 22(4):215–233, 2010.

[48] Amos Korman, David Peleg, and Yoav Rodeh. Constructing labeling schemes through universal matrices. *Algorithmica*, 57(4):641–652, 2010.

Publications in Conferences and Workshops

[49] Olivier Bailleux, Yacine Boufkhad, and Olivier Roussel. New encodings of pseudo-boolean constraints into cnf. In *Theory and Applications of Satisfiability Testing - SAT 2009*, pages 181–194, Royaume-Uni Swansea, 2009.

[50] Hervé Baumann, Pierluigi Crescenzi, and Pierre Fraigniaud. Parsimonious flooding in dynamic graphs. In *Proc. 28th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 260–269, 2009.

[51] Hervé Baumann and Pierre Fraigniaud. Sub-linear universal spatial gossip protocols. In *Proc. 16th International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, LNCS. Springer, 2009.

[52] Farid Benbadis, Nidhi Hegde, Fabien Mathieu, and Diego Perino. Playing with the bandwidth conservation law. In *International Conference on Peer-to-Peer Computing (P2P)*, 2008.

[53] Thomas Bonald, Laurent Massoulie, Fabien Mathieu, Diego Perino, and Andy Twigg. Epidemic live streaming: Optimal performance trade-offs. In *International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, 2008.

[54] Y. Boufkhad, F. Mathieu, F. de Montgolfier, D. Perino, and L. Viennot. Achievable catalog size in peer-to-peer video-on-demand systems. In *Proc. of the 7th Int. Workshop on Peer-to-Peer Systems (IPTPS)*, pages 1–6, 2008.

[55] Yacine Boufkhad and Thomas Hugel. Non uniform selection of solutions for upper bounding the 3-sat threshold. In *SAT*, pages 99–112, 2010.

[56] Yacine Boufkhad, Fabien Mathieu, Fabien de Montgolfier, Diego Perino, and Laurent Viennot. Fine tuning of a distributed vod system. In *18th IEEE International Conference on Computer Communications and Networks (ICCCN)*, pages 1–7, 2009.

[57] Yacine Boufkhad, Fabien Mathieu, Fabien de Montgolfier, Diego Perino, and Laurent Viennot. An upload bandwidth threshold for peer-to-peer video-on-demand scalability. In *23rd IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 1–10, 2009.

[58] A. Chaintreau, P. Fraigniaud, and E. Lebhar. Networks become navigable as nodes move and forget. In *35th International Colloquium on Automata, Languages and Programming (ICALP)*, LNCS 5125, pages 133–144. Springer, 2008.

[59] A. Chaintreau, P. Fraigniaud, and E. Lebhar. Opportunistic spatial gossip over mobile social networks. In *1st ACM Workshop on Online Social Networks (WOSN)*, pages 73–78, 2008.

[60] Julien Clément, Carole Delporte-Gallet, Hugues Fauconnier, and Mihaela Sighireanu. Guidelines for the verification of population protocols. In *ICDCS*, pages 215–224, 2011.

[61] Julien Clément, Carole Delporte-Gallet, Hugues Fauconnier, and Mihaela Sighireanu. Mode d'emploi pour la vérification des protocoles de population. In Pascal Ducourthial, Bertrand et Felber, editor, *13es Rencontres Francophones sur les Aspects Algorithmiques de Télécommunications (AlgoTel)*, Cap Estérel, France, 2011.

[62] Pierluigi Crescenzi, Roberto Grossi, Michel Habib, Lorenzo Lanzi, and Andrea Marino. On computing the diameter of real-world undirected graphs. In *Workshop ESA in honor of Giorgio Ausiello, submitted to TCS*, 2011.

[63] Fabien de Montgolfier, Mauricio Soto, and Laurent Viennot. Asymptotic modularity of some graph classes. In *22nd International Symposium on Algorithms and Computation (ISAAC)*, pages 435–444, Yokohama, 2011.

[64] Fabien de Montgolfier, Mauricio Soto, and Laurent Viennot. Clustering de métrique et clustering de graphe. In Pascal Ducourthial, Bertrand et Felber, editor, *13es Rencontres Francophones sur les Aspects Algorithmiques de Télécommunications (AlgoTel)*, Cap Estérel, France, 2011. Equipe-projet GANG (inria).

[65] Fabien de Montgolfier, Mauricio Soto, and Laurent Viennot. Treewidth and hyperbolicity of the internet. In *10th IEEE International Symposium on Network Computing and Applications (IEEE NCA)*, pages 1–8, Boston, 2011.

[66] Carole Delporte-Gallet, Stéphane Devismes, and Hugues Fauconnier. Approximation of *elta*-timeliness. In *SSS*, pages 435–451, 2010.

[67] Carole Delporte-Gallet, Stéphane Devismes, Hugues Fauconnier, and Mikel Larrea. Algorithms for extracting timeliness graphs. In *SIROCCO*, pages 127–141, 2010.

[68] Carole Delporte-Gallet, Hugues Fauconnier, Eli Gafni, and Petr Kuznetsov. Brief announcement: On the meaning of solving a task with a failure detector. In *DISC*, pages 145–146, 2011.

[69] Carole Delporte-Gallet, Hugues Fauconnier, Rachid Guerraoui, and Anne-Marie Kermarrec. Brief announcement: byzantine agreement with homonyms. In *SPAA*, pages 74–75, 2010.

[70] Carole Delporte-Gallet, Hugues Fauconnier, Rachid Guerraoui, Anne-Marie Kermarrec, Eric Ruppert, and Hung Tran-The. Byzantine agreement with homonyms. In *30th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 21–30, 2011.

[71] B. Derbel, C. Gavoille, D. Peleg, and L. Viennot. On the locality of distributed sparse spanner construction. In ACM Press, editor, *27th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 273–282, 2008.

[72] Bilel Derbel, Cyril Gavoille, David Peleg, and Laurent Viennot. Local computation of nearly additive spanners. In *DISC*, volume 5805, pages 176–190, Espagne Elche/Elx, 2009. Springer.

[73] Yuval Emek, Pierre Fraigniaud, Amos Korman, and A. Rosen. On the additive constant of the k-server work function algorithm. In *7th Workshop on Approximation and Online Algorithms (WAOA)*, LNCS. Springer, 2009.

[74] Yuval Emek, Pierre Fraigniaud, Amos Korman, and A. Rosen. Online computation with advice. In *Proc. 36th International Colloquium on Automata, Languages and Programming (ICALP)*, LNCS 5555, pages 427–438, 2009.

[75] Yuval Emek and Amos Korman. New bounds for the controller problem. In *23rd International Symposium on Distributed Computing (DISC)*, volume 5805 of *Lecture Notes in Computer Science*, pages 22–34, 2009.

[76] Yuval Emek and Amos Korman. Efficient threshold detection in a distributed environment: extended abstract. In *29th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, Lecture Notes in Computer Science, pages 183–191, 2010.

[77] Yuval Emek, Amos Korman, and Yuval Shavitt. Approximating the statistics of various properties in randomly weighted graphs. In *22nd Annual ACM-SIAM Symposium on Discrete Algorithms, (SODA)*, pages 1455–1467, San Francisco, 2011.

[78] Pierre Fraigniaud and George Giakkoupis. The effect of power-law degrees on the navigability of small worlds. In *Proc. 28th ACM Symposium on Principles of Distributed Computing (PODC)*, pages 240–249, 2009.

[79] Pierre Fraigniaud and George Giakkoupis. On the bit communication complexity of randomized rumor spreading. In *SPAA*, pages 134–143, 2010.

[80] Pierre Fraigniaud and George Giakkoupis. On the searchability of small-world networks with arbitrary underlying structure. In *STOC*, pages 389–398, 2010.

[81] Pierre Fraigniaud and Amos Korman. On randomized representations of graphs using short labels. In *21st Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 131–137, 2009.

[82] Pierre Fraigniaud and Amos Korman. Compact ancestry labeling schemes for xml trees. In *21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, Lecture Notes in Computer Science, pages 458–466, 2010.

[83] Pierre Fraigniaud and Amos Korman. An optimal ancestry scheme and small universal posets. In *42nd ACM Symposium on Theory of Computing (STOC)*, Lecture Notes in Computer Science, pages 611–620, 2010.

[84] Pierre Fraigniaud, Amos Korman, and David Peleg. Local distributed decision. In *52nd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, Palm Springs, 2011.

[85] Pierre Fraigniaud, Emmanuelle Lebhar, and Laurent Viennot. The inframetric model for the internet. In *Proceedings of the 27th IEEE International Conference on Computer Communications (INFOCOM)*, pages 1085–1093, Phoenix, 2008.

[86] Pierre Fraigniaud and Andrzej Pelc. Delays induce an exponential memory gap for rendezvous in trees. In *SPAA*, pages 224–232, 2010.

[87] Pierre Fraigniaud and Andrzej Pelc. Delays induce an exponential memory gap for rendezvous in trees. In *10th Latin American Theoretical Informatics Symposium (LATIN)*. Springer, 2012.

[88] Pierre Fraigniaud and Hoang Anh Phan. Degree hunter: on the impact of balancing node degrees in de bruijn-based overlay networks. In *24th IEEE International Parallel and Distributed Processing Symposium, Workshops and Phd Forum (IPDPSW)*, 2010.

[89] Pierre Fraigniaud and Hoang Anh Phan. "tree-farms" for tree-based multicast schemes in peer-to-peer overlay networks. In *2010 IEEE International Conference In Communications (ICC)*, 2010.

[90] Pierre Fraigniaud, Sergio Rajsbaum, and Corentin Travers. Locality and checkability in wait-free computing. In *25th International Symposium on Distributed Computing (DISC)*, volume 6950 of *LNCS*, pages 333–347. Springer, 2011.

[91] Cyril Gavoille, Quentin Godfroy, and Laurent Viennot. Graphes de recouvrement multichemins. In *12èmes Rencontres Francophones sur les Aspects Algorithmiques de Télécommunications (AlgoTel)*, France Belle Dune, 2010. Maria Gradinariu Potop-Butucaru and Hervé Rivano.

[92] Cyril Gavoille, Quentin Godfroy, and Laurent Viennot. Multipath spanners. In *Structural Information and Communication Complexity, 17th International Colloquium (SIROCCO)*, pages 211–223, Sirince, 2010.

[93] Cyril Gavoille, Quentin Godfroy, and Laurent Viennot. Node-disjoint multipath spanners and their relationship with fault-tolerant spanners. In *15th International Conference on Principles of Distributed Systems (OPODIS)*, pages 1–16, Toulouse, 2011.

[94] Michel Habib and Juraj Stacho. Unique perfect phylogeny is *np*-hard. In *CPM*, pages 132–146, 2011.

[95] Michel Habib and Thu-Hien To. On a conjecture about compatibility of multi-states characters. In *WABI*, pages 116–127, 2011.

[96] Nidhi Hegde, Fabien Mathieu, and Diego Perino. Diffusion épidémique de chunks en quasi-direct : la taille compte. In Augustin Chaintreau and Clemence Magnien, editors, *Algotel*, France Carry-Le-Rouet, 2009.

[97] Philippe Jacquet and Laurent Viennot. Average size of unstretched remote-spanners. In *5th ACM Workshop on Analytic Algorithmics and Combinatorics (ANALCO)*, pages 23–33, 2009.

[98] Philippe Jacquet and Laurent Viennot. Remote spanners: what to know beyond neighbors. In *23rd IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 1–10, 2009.

[99] Liah Kor, Amos Korman, and David Peleg. Tight bounds for distributed mst verification. In *28th International Symposium on Theoretical Aspects of Computer Science, (STACS)*, pages 69–80, Dortmund, 2011.

[100] Amos Korman. Compact routing schemes for dynamic trees in the fixed port model. In *10th International Conference on Distributed Computing and Networking, (ICDCN)*, volume 5408 of *Lecture Notes in Computer Science*, pages 218–229, 2009.

[101] Amos Korman, Shay Kutten, and Toshimitsu Masuzawa. Fast and compact self stabilizing verification, computation, and fault detection of an mst. In ACM Press, editor, *30th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 311–320, San Jose, 2011.

[102] Amos Korman, Jean-Sébastien Sereni, and Laurent Viennot. Toward more localized local algorithms: Removing assumptions concerning global knowledge. In ACM Press, editor, *30th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 49–58, San Jose, 2011.

[103] Fabien Mathieu. Live seeding: Performance bounds of seeders for P2P live streaming. In *P2P '11: Proceedings of the 2011 Eleventh International Conference on Peer-to-Peer Computing*, pages 172–181. IEEE Computer Society, 2011.

[104] Fabien Mathieu and Laurent Viennot. Local aspects of the global ranking of web pages. In *10th International Conference on Innovative Internet Community Services (I2CS)*, pages 493–506, Neuchâtel, 2010.

[105] Luca Muscariello and Diego Perino. Evaluating the performance of multi-path routing and congestion control in presence of network resource management. In *IEEE Conference on Ultra Modern Telecommunications (ICUMT)*, 2009.

[106] Luca Muscariello and Diego Perino. Modeling multi-path routing and congestion control under fifo and fair queuing. In *IEEE Conference on Local Computer Network (LCN)*, 2009.

[107] Luca Muscariello, Diego Perino, and Bruno Nardelli. Towards real implementations of dynamic robust routing paradigms exploiting path diversity. In *IEEE Conference on Ultra Modern Telecommunications (ICUMT)*, 2009.

[108] Luca Muscariello, Diego Perino, and Dario Rossi. Do next generation networks need path diversity? In *International Conference on Communications (ICC)*, 2009.

[109] Atish Das Sarma, Stephan Holzer, Liah Kor, Amos Korman, Danupon Nanongkai, Gopal Pandurangan, David Peleg, and Roger Wattenhofer. Distributed verification and hardness of distributed approximation. In ACM Press, editor, *43rd ACM Symposium on Theory of Computing, (STOC)*, pages 363–372, San Jose, 2011.

General audience papers

[110] Fabien de Montgolfier and Vincent Limouzy. Stratégies de capture de fugitifs... ou l'application de la théorie des graphes à pacman.)I(nterstices, 2009.

[111] Laurent Viennot. Idée reçue : Web et Internet, c'est la même chose.)I(nterstices, 2009.

[112] Laurent Viennot. Les réseaux ont la fibre de l'information.)I(nterstices, 2009.

Internal Reports

[113] François Baccelli, Fabien Mathieu, and Ilkka Norros. Performance of P2P Networks with Spatial Interactions of Peers. Rapport de recherche RR-7713, INRIA, August 2011.

[114] Yacine Boufkhad and Thomas Hugel. Estimating satisfiability. 2010.

[115] Cyril Gavoille, Quentin Godfroy, and Laurent Viennot. Node-Disjoint Multipath Spanners and their Relationship with Fault-Tolerant Spanners. Technical report, LaBRI, September 2011. 15 pages.

[116] Michel Habib and Juraj Stacho. Unique perfect phylogeny is np-hard. 2010.

[117] Nidhi Hegde, Fabien Mathieu, and Diego Perino. Size does matter (in P2P live streaming). Rapport de recherche, INRIA, 2009. RR-7032.

[118] Fabien Mathieu. On using seeders for P2P live streaming. Technical Report RR-7608, INRIA, 2011.

[119] Fabien Mathieu and Diego Perino. On resource aware algorithms in epidemic live streaming. Rapport de recherche, INRIA, 2009. RR-7031.

[120] Andreas Neuenkirch, Samy Tindel, and Jérémie Unterberger. Discretizing the fractional levy area. 2009.