INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

# *Project-Team Gang*

# *Graphs, networks and algorithms*

*Paris - Rocquencourt*

Theme : Networks and Telecommunications

# Table of contents


**1. Team** . . . . . 1
**2. Overall Objectives** . . . . . 1
**3. Scientific Foundations** . . . . . 2
3.1.1. Structured/Unstructured Overlays 2
3.1.2. Small World Phenomenon 2
3.1.3. Doubling Metrics 2
3.1.4. Bounded Width Classes of Graphs 3
**4. Application Domains** . . . . . 3
**5. New Results** . . . . . 3
5.1. Small world networks structure 3
5.1.1. Small world routing 4
5.1.2. Rumor spreading 4
5.2. Peer-to-Peer 5
5.2.1. Video-on-Demand over Peer-to-Peer networks 5
5.2.2. Unstructured P2P live streaming 5
5.3. Dynamic graph algorithms 6
5.3.1. Routing: design, modeling and analysis 6
5.3.1.1. Multi-path routing 6
5.3.1.2. Compact Routing Schemes for Dynamic Trees in the Fixed Port Model 6
5.3.1.3. Adversarial Queuing Theory with Setups 7
5.3.2. Online algorithms 7
5.3.2.1. New Bounds for the Controller Problem 7
5.3.2.2. Online Computation with Advice 7
5.4. Understanding graph representations 8
5.4.1. Graph labeling 8
5.4.1.1. Informative labeling schemes for trees 8
5.4.1.1.1. Ancestry labeling 8
5.4.1.1.2. On randomized representations of graphs using short labels 8
5.4.1.2. Compact labeling for spanning trees 9
5.4.1.3. Labeling schemes for trees with queries 9
5.4.2. Efficient graph spanners 10
5.4.2.1. Distributed Computation of Sparse Spanner 10
5.4.2.2. Spanners for Ad Hoc Networks 10
5.4.3. Graph decompositions 11
5.4.3.1. NLC decomposition 11
5.4.3.2. Umodular decomposition 11
5.5. Discrete Optimization Algorithms 11
5.5.1. Efficient encoding of pseudo-boolean constraints 11
5.5.2. Three way decomposition of permutation problems 11
**6. Contracts and Grants with Industry** . . . . . 12
6.1.1. Measuring Internet with and for peer-to-peer networks 12
6.1.2. Modeling node and connection dynamics 12
6.1.3. Peer-to-peer application design 12
**7. Other Grants and Activities** . . . . . 12
7.1. National initiatives 12
7.2. Actions Funded by the EC 13
**8. Dissemination** . . . . . 13
8.1. Services to the scientific community 13
8.2. Teaching 13

8.3. Theses 14
8.3.1. Defended theses 14
8.3.2. Ongoing theses 14
**9. Bibliography** . . . . . **14**

*Gang is a joined team between INRIA, CNRS and Paris Diderot University, through the "laboratoire d'informatique algorithmique, fondements et applications",* LIAFA *(UMR 7089).*

# 1. Team

**Research Scientist**

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

**Faculty Member**

Michel Habib [ Professor (Paris Diderot Univ.), HdR ]
Yacine Boufkhad [ Assistant professor (Paris Diderot Univ.) ]
Pierre Charbit [ Assistant professor (Paris Diderot Univ.) ]
Fabien de Montgolfier [ Assistant professor (Paris Diderot Univ.) ]

**PhD Student**

Damien Noguës [ PhD Student (AMN) ]
Diego Perino [ PhD Student (CIFRE FranceTelecom) ]
Anh Hoand Phan [ PhD Student (BDI) ]
Mauricio Soto [ PhD Student (Chile) ]
Hien Hu To [ PhD Student (AMX) ]

**Administrative Assistant**

Christine Anocq [ shared time (with Hipercom) ]

# 2. Overall Objectives

## 2.1. Overall Objectives

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

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

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

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

# 3. Scientific Foundations

## 3.1. Scientific Foundations

### 3.1.1. Structured/Unstructured Overlays

Recent years have brought tremendous progress along the peer-to-peer paradigm allowing large scale decentralized application to be practically deployed. The main achievement of this trend is certainly efficient content distribution through the BitTorrent protocol [34]. The power of peer-to-peer content distribution is to rely on the upload capacity of the node interested in receiving the content. This allows to scale to very large number of participants. The main breakthrough of BitTorrent resides in its "tit for tat" strategy inspired from game theory to give incentive to cooperate. For that purpose, a peer preferentially uploads preferentially offering best reciprocity. This kind of preferences induces an interesting graph structure with ordered neighborhoods. Understanding the dynamic behavior of such affinity graphs is an important for stabilizing the performance of such protocols.

A second major achievement of the peer-to-peer paradigm concerns indexing with distributed hashtables [50], [52], [54]. The idea behind these proposals is to organize the peers into a structure close to well known graphs with low diameter such as hight dimension torus, hypercube or de Bruijn graph. Efficient routing to the node storing a given key is then guaranteed. This academic work has lead to practical basic indexing facilities by introducing redundancy in the structure [49]. This is typically the kind of approach we want to promote: from known efficient theoretical solutions to practical working protocols. We have contributed to this trend by introducing de Bruijn based solutions [42], [43] with redundancy in the contact graph to resist node churn.

### 3.1.2. Small World Phenomenon

Popularized emerging properties include degree distributions observed to be power law in many networks or clustering coefficient observed to be high in social networks or low average distances. This last point gave the denomination of "small worlds" for this type of networks. Some work [31], [55] try to give models that give raise to such statistical properties. In that line, numerous results such as [30] try to derive efficient algorithms based only on these statistical properties. This particular approach tends to concentrate load on nodes with high degrees and may not be suited for applications where nodes have similar capacities. Other interesting work [39] try to explain this statistical observation forms an inherent optimization problem operating when constructing the network.

On the other hand, in its seminal paper [46], Kleinberg focuses on the algorithmic aspects of such social networks and shows how adding random links to a torus can produce efficient greedy routing. This result has been extend to more general classes of graphs [53], [40] such as bounded growth metric graphs [38]. However, this augmentation process is not always feasible [41]. Such theoretical work is particularly interesting for overlay networks where this augmenting process simply consists in opening additional connections.

### 3.1.3. Doubling Metrics

Bounded growth and doubling metrics generalize Euclidean metrics. A metric has bounded growth if the size of any ball increases by a factor not larger than $2^d$ when its radius is doubled [45]. A metric is doubling if any set of diameter $D$ can be covered with at most $2^d$ sets of diameter $D/2$ [32]. In both cases, the smallest acceptable value of $d$ is called the dimension of the metric. The metric of any $d$ dimensional Euclidean space has bounded growth dimension $O(d)$. Any bounded growth metric of dimension $d$ has doubling dimension $O(d)$. The doubling metric is the most general and has the additional property of being inherited by subspaces: the metric induced by a doubling metric on a subset of nodes is also doubling. For example, sampling nodes in an Euclidean space always results in a doubling metric.

As networks are embedded in our usual three dimensional space, it is legitimate to think than some network metrics may be modeled through doubling metrics. Recent results thus investigate network problems for the restricted classes of graphs with bounded growth or doubling metric [53], [29], [45], [47]. However, the doubling nature of large networks such as the Internet has still not be tested.

### *3.1.4. Bounded Width Classes of Graphs*

Many graph parameters such as treewidth, branchwidth, rankwidth, cutwidth, cliquewidth ...have been introduced in recent years [51], [36] to measure the structure of a given graph. These parameters are of course NP-complete to compute, but when it can be proved that for a given class of graphs the parameter is bounded by a constant then it can be proved using graph grammars (see Courcelle's fundamental work) that most of the optimization problems on this class are polynomially tractable, and sometimes we know the existence of a linear algorithm (but the hidden constant can be very high !)

The most famous parameter, namely the treewidth captures the distance of the graph to a tree, and therefore when the treewidth is small a dynamic programming approach can be used [35].

Despite some promising results, applications of these notions has still to be done for networks, in a practical perspective.

# 4. Application Domains

## 4.1. Application Domains

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

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

# 5. New Results

## 5.1. Small world networks structure

**Participants:** Pierre Fraigniaud, Cyril Gavoille [CNRS LABRI, University of Bordeaux, France], George Giakkoupis [CNRS LIAFA, University of Paris Diderot, France], Adrian Kosowski [CNRS LABRI, University of Bordeaux, France], Zvi Lotker [University of Tel Aviv, Israel].

### 5.1.1. Small world routing

In [12], we analyze decentralized routing in small-world networks that combine a wide variation in node degrees with a notion of spatial embedding. Specifically, we consider a variation of Kleinberg's augmented-lattice model (STOC 2000), where the number of long-range contacts for each node is drawn from a power-law distribution. This model is motivated by the experimental observation that many "real-world" networks have power-law degrees. In such networks, the exponent $\alpha$ of the power law is typically between 2 and 3. We prove that, in our model, for this range of values, $2 < \alpha < 3$, the expected number of steps of greedy routing from any source to any target is $O(\log^{\alpha-1} n)$ steps. This bound is tight in a strong sense. Indeed, we prove that the expected number of steps of greedy routing for a uniformly-random pair of source–target nodes is $\Omega(\log^{\alpha-1} n)$ steps. We also show that, for $\alpha < 2$ or $\alpha \geq 3$, greedy routing performs in $\Theta(\log^2 n)$ expected steps, and for $\alpha = 2$, $\Theta(\log^{1+\epsilon} n)$ expected steps are required, where $1/3 \leq \epsilon \leq 1/2$. To the best of our knowledge, these results are the first to formally quantifying the effect of the power-law degree distribution on the navigability of small worlds. Moreover, they show that this effect is significant. In particular, when $\alpha$ approaches 2 from above, the expected number of steps of greedy routing in the augmented lattice with *power-law degrees* approaches the square-root of the expected number of steps of greedy routing in the augmented lattice with *fixed degrees*, although both networks have the same *average degree*.

### 5.1.2. Rumor spreading

*Gossip* protocols are communication protocols in which, periodically, every node of a network exchanges information with some other node chosen according to some (randomized) strategy. These protocols have recently found various types of applications for the management of distributed systems. *Spatial* gossip protocols are gossip protocols that use the underlying spatial structure of the network, in particular for achieving the "closest-first" property. This latter property states that the closer a node is to the source of a message the more likely it is to receive this message within a prescribed amount of time. Spatial gossip protocols find many applications, including the propagation of alarms in sensor networks, and the location of resources in P2P networks. In [6], we design a sub-linear spatial gossip protocol for arbitrary graphs metric. More specifically, we prove that, for any graph metric with maximum degree $\Delta$, for any source $s$ and any ball centered at $s$ with size $b$, new information is spread from $s$ to all nodes in the ball within $O((\sqrt{b \log b} \log\log b + \Delta) \log b)$ rounds, with high probability. Moreover, when applied to general metrics with uniform density, the same protocol achieves a propagation time of $O(\log^2 b \log\log b)$ rounds.

An edge-Markovian process with birth-rate $p$ and death-rate $q$ generates sequences of graphs $(G_0, G_1, G_2, ...)$ 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 $q$. Clementi et al. (PODC 2008) analyzed thoroughly information dissemination in such dynamic graphs, by establishing bounds on their flooding time — 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 $t' > t$. In [5], we establish tight bounds on the complexity of flooding for all possible birth rates and death rates, completing the previous results by Clementi et al. Moreover, we note that despite its many advantages in term of simplicity and robustness, flooding suffers from its high bandwidth consumption. Hence we al! so show that flooding in dynamic graphs can be implemented in a more 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 sa! me 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.

## 5.2. Peer-to-Peer

### *5.2.1. Video-on-Demand over Peer-to-Peer networks*

**Participants:** Yacine Boufkhad, Fabien Mathieu [Orange Labs, Issy Les Moulineaux, France], Fabien de Montgolfier, Diego Perino, Laurent Viennot.

We consider the fully distributed Video-on-Demand problem, where $n$ nodes called *boxes* store a large set of videos and collaborate to serve simultaneously $n$ videos or less between them. It is said to be *scalable* when $\Omega(n)$ videos can be distributively stored under the condition that any sequence of demands for these videos can always be satisfied. Our main result consists in establishing a threshold on the average upload bandwidth of a box, above which the system becomes scalable. We are thus interested in the normalized upload capacity $u = \frac{upload\ bandwidth}{video\ bitrate}$ of a box. The number $m$ of distinct videos stored in the system is called its catalog size.

In [7], we show an upload capacity threshold of 1 for scalability in a homogeneous system, where all boxes have the same upload capacity. More precisely, a system with $u < 1$ has constant catalog size $m = O(1)$ (every box must store some data of every video). On the other hand, for $u > 1$, an homogeneous system where all boxes have same upload capacity at least $u$ admits a static allocation of $m = \Omega(n)$ videos into the boxes such that any adversarial sequence of video demands can be satisfied. Moreover, such an allocation can be obtained randomly with high probability. This result is generalized to a system of boxes that have heterogeneous upload capacities under some balancing conditions.

In [8], by means of extensive simulations we analyze the impact of: i) the video allocation technique used for distributed storage ii) the use of cache to allow nodes to re-distribute the video they are downloading iii) the use of static/dynamic algorithms for video distribution.

Based on these results, we provide some guidelines for setting the system parameters: the use of cache strongly improves system performance; popularity based allocation techniques can be sensitive and bring little improvement; dynamic distribution algorithms are needed only in extreme scenarios while static ones are commonly sufficient.

### *5.2.2. Unstructured P2P live streaming*

**Participants:** Nidhi Hegde [Orange Labs, Issy Les Moulineaux, France], Fabien Mathieu [Orange Labs, Issy Les Moulineaux, France], Diego Perino [Orange Labs, Issy Les Moulineaux, France].

In unstructured P2P live streaming systems the stream is not forwarded as a continuous flow of data but is divided in a series of pieces (chunks), that are injected in the system by a source and exchanged among peers in order to retrieve the complete sequence and play out the stream. Data exchange is therefore driven by chunk exchange algorithms run locally by nodes, which can be described by their chunk/peer selection policies. It turns out that the most popular commercial peer-to-peer systems for live streaming like CoolStreaming, PPLive, SopCast, are based on such an unstructured approach.

The performance trade-offs of chunk exchange algorithms have been deeply analyzed for homogeneous systems, where all peers have the same upload capacity. In [27] we analyze chunk diffusion algorithms designed for heterogeneous environments, where peers have different upload capacities. We focus on the peer selection process and propose a generic model that encompasses a large class of algorithms. We derive recursive formulas to describe the chunk diffusion function of a generic *latest blind chunk / resource aware peer* selection scheme. By means of simulations, we analyze the resource awareness-agnostism trade-offs on the peer selection process and the impact of the source distribution policy in non-homogeneous networks. We highlight that the early diffusion of a given chunk is crucial for its overall diffusion performance, and a fairness trade-off arises between the performance of heterogeneous peers, as a function of the level of awareness. Moreover, we show the critical role the source selection policy plays on chunk diffusion performance.

A good diffusion scheme is indeed essential for the performance of an unstructured P2P live streaming system. For a given scheme however, an optimization at a detailed level is also important. This involves the fine tuning of dissemination parameters, such as chunk size, receiver buffer size, number of peers to probe, etc. In [26], [23] we investigate optimal sizing of chunks and probe sets, i.e. the number peers a given node probes before transmitting chunks. The analysis is performed by means of an event-based simulator. We show that the chunk size significantly impacts performance and that it should fall within a given range which is mostly determined by the median RTT of the network and the stream rate. We also show that the size of the probe set affects performance of diffusion schemes, and, in particular, a probe set larger than the actual number of concurrent connections may improve miss ratio/delay performance by modifying the suitable chunk size ranges.

## 5.3. Dynamic graph algorithms

### *5.3.1. Routing: design, modeling and analysis*

*5.3.1.1. Multi-path routing*

**Participants:** Luca Muscariello [Orange Labs, Issy Les Moulineaux, France], Bruno Nardelli [Rice University, Texas, USA], Diego Perino [Orange Labs, Issy Les Moulineaux, France], Dario Rossi [Telecom ParisTech, France].

Network traffic is increasing in size and is becoming more and more dynamic leading to unpredictable and highly variable traffic patters. Multi-path routing is considered a powerful approach to deal with the aforementioned traffic patterns, mostly for intra-domain TE, inter-AS path selection under the same ISP, routing in wireless networks and metropolitan access networks. Through an optimization framework in [22] we quantify the benefit in using multi-path routing and we analyze the role network topology and traffic matrix play on multi-path routing performance. Starting from the insights inferred from this analysis in [20], [19], we introduce MIRTO, which is a distributed multi-path routing protocol that jointly uses best path selection and flow control for optimality and stability. In [20], [19] we introduce an analytical model in order to compare different routing schemes using fluid ordinary differential equations (ODEs). We model the sending rate of MIRTO and of other two recently proposed algorithms, TEXCP and TRUMP, and we compare their performance on the Abilene network topology with FIFO and Fair-Queuing scheduling. We show MIRTO consumes less network resource but suffers of longer convergence times because it relies on simpler feedbacks and it is not equation based. We also highlight that Fair-Queuing scheduling leads to sub-optimal global resource allocations because it imposes fairness at link level. We implement MIRTO, TEXCP and TRUMP in a software prototype in order to analyze their performance in real networks under real traffic conditions. Results reported in [21] confirm the trends already highlighted by our analytical analysis, and show that dynamic yet stable traffic engineering is not only feasible but expected with increasing interest by network operators.

*5.3.1.2. Compact Routing Schemes for Dynamic Trees in the Fixed Port Model*

**Participant:** Amos Korman.

This paper considers the routing problem in dynamic trees under the fixed-port model, in which an adversary chooses the port numbers assigned to each node. We present two routing schemes for dynamic trees that maintain labels of asymptotically optimal size using extremely low average message complexity (per node). Specifically, we first present a dynamic routing scheme that supports additions of both leaves and internal nodes, maintains asymptotically optimal labels and incurs only $O(\log^2 n/\log^2 \log n)$ average message complexity. This routing scheme is then extended to supports also deletions of nodes of degree at most 2. The extended scheme incurs $O(\log^2 n)$ average message complexity and still maintains asymptotically optimal labels.

We would like to point out that the best known routing scheme for dynamic trees that maintains asymptotically optimal labels in the fixed port model has very high average message complexity, namely, $O(n^\epsilon)$. Moreover, that scheme supports additions and removals of leaf nodes only.

*5.3.1.3. Adversarial Queuing Theory with Setups*

**Participant:** Mauricio Soto.

We look at routing and scheduling problems on Kelly type networks [3] where the injection process is under the control of an adversary. The novelty of the model we consider is that the adversary injects requests of distinct types. Resources are subject to switch-over delays or setups when they begin servicing a new request class. In this new setting, we study the behavior of sensible policies as introduced by Dai and Jennings [37].

We first show that the model is robust in the sense that under some mild conditions universal stability of work conserving packet routing protocols is preserved for natural variants of the underlying model. Also, the model's equivalence to so called token networks is established.

We adapt, to the multi-type request and setup setting, standard arguments for proving stability. Nevertheless, we provide counterexamples that show that for several reasonable adaptations of contention resolution protocols to the multi-type case, stability results do not carry over from the single-type scenario. This motivates us to explore fluid model based arguments that could be used for proving stability for a given network. Specifically we show analogues of results obtained by Gamarnik [44] but in the multi-type request with setups scenario.

## *5.3.2. Online algorithms*

*5.3.2.1. New Bounds for the Controller Problem*

**Participants:** Yuval Emek [University of Tel Aviv, Israel], Amos Korman.

The $(M, W)$-controller, originally studied by Afek, Awerbuch, Plotkin, and Saks, is a basic distributed tool that provides an abstraction for managing the consumption of a global resource in a distributed dynamic network. The input to the controller arrives online in the form of requests presented at arbitrary nodes. A request presented at node $u$ corresponds to the "desire" of some entity to consume one unit of the global resource at $u$ and the controller should handle this request within finite time by either granting it with a permit or denying it. Initially, $M$ permits (corresponding to $M$ units of the global resource) are stored at a designated root node. Throughout the execution permits can be transported from place to place along the network's links so that they can be granted to requests presented at various nodes; when a permit is granted to some request, it is eliminated from the network. The fundamental rule of an $(M, W)$-controller is that a request should not be denied unless it is certain that at least $M - W$ permits are eventually granted. The most efficient $(M, W)$-controller known to date has message complexity $O(N \log^2 N \log \frac{M}{W+1})$, where $N$ is the number of nodes that ever existed in the network (the dynamic network may undergo node insertions and deletions).

In this paper we establish two new lower bounds on the message complexity of the controller problem. We first prove a simple lower bound stating that any $(M, W)$-controller must send $\Omega(N \log \frac{M}{W+1})$ messages. Second, for the important case when $W$ is proportional to $M$ (this is the common case in most applications), we use a surprising reduction from the (centralized) monotonic labeling problem to show that any $(M, W)$-controller must send $\Omega(N \log N)$ messages. In fact, under a long lasting conjecture regarding the complexity of the monotonic labeling problem, this lower bound is improved to a tight $\Omega(N \log^2 N)$. The proof of this lower bound requires that $N = O(M)$ which turns out to be somewhat inevitable due to a new construction of an $(M, M/2)$-controller with message complexity $O(N \log^2 M)$.

*5.3.2.2. Online Computation with Advice*

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

In [11], we consider a model for online computation in which the online algorithm receives, together with each request, some information regarding the future, referred to as *advice*. The advice provided to the online algorithm may allow an improvement in its performance, compared to the classical model of complete lack of information regarding the future. We are interested in the impact of such advice on the competitive ratio, and in particular, in the relation between the size $b$ of the advice, measured in terms of bits of information per request, and the (improved) competitive ratio. Since $b = 0$ corresponds to the classical online model,

and $b = \lceil \log |\mathcal{A}| \rceil$, where $\mathcal{A}$ is the algorithm's action space, corresponds to the optimal (offline) one, our model spans a spectrum of settings ranging from classical online algorithms to offline ones. We illustrate the applicability of this model by considering two of the most extensively studied online problems, namely, *metrical task systems (MTS)* and the *k-server* problem. For MTS we establish tight (up to constant factors) upper and lower bounds on the competitive ratio of deterministic and randomized online algorithms with advice for any choice of $1 \leq b \leq \Theta(\log n)$, where $n$ is the number of states in the system: we prove that any randomized online algorithm for MTS has competitive ratio $\Omega(\log(n)/b)$ and we present a deterministic online algorithm for MTS with competitive ratio $O(\log(n)/b)$. For the $k$-server problem we construct a deterministic online algorithm for general metric spaces with competitive ratio $k^{O(1/b)}$ for any choice of $\Theta(1) \leq b \leq \log k$.

In [10], we consider the Work Function Algorithm for the k-server problem. We show that if the Work Function Algorithm is c-competitive, then it is also strictly (2c)-competitive. As a consequence of [Koutsoupias and Papadimitriou, JACM 1995] this also shows that the Work Function Algorithm is strictly (4k-2)-competitive.

## 5.4. Understanding graph representations

### *5.4.1. Graph labeling*

*5.4.1.1. Informative labeling schemes for trees*

**Participants:** Pierre Fraigniaud, Amos Korman.

5.4.1.1.1. Ancestry labeling

In [14], we consider *ancestry labeling schemes*. Such a scheme labels the nodes of any tree in such a way that ancestry queries between any two nodes can be answered just by looking at their corresponding labels. The common measure to evaluate the quality of an ancestry scheme is by its *label size*, that is the maximum number of bits stored in a label, taken over all $n$-node trees. The design of ancestry labeling schemes finds applications in XML search engines. In these contexts, even small improvements in the label size are important. In fact, the literature about this topic is interested in the exact label size rather than just its order of magnitude. As a result, following the proposal of a simple interval based ancestry scheme with label size $2 \log n$ bits (Kannan et al., SIDMA 92), a considerable amount of work was devoted to improve the bound on the label size. The current state of the art upper bound is $\log n + O(\sqrt{\log n})$ bits (Abiteboul et al., SICOMP 06) which is still far from the known $\log n + \Omega(\log \log n)$ lower bound (Alstrup et al., SODA 03). Moreover, the hidden constant factor in the additive $O(\sqrt{\log n})$ term is large, which makes that scheme less efficient than the simple interval based scheme for typical current XML trees. It is for this reason that Kaplan et al. suggested other ancestry schemes which, on average, use labels whose size, for typical XML data, is smaller by about 10%-30% than the ones used for the interval based scheme (SODA 02). Motivated by the fact that typical XML trees have extremely small depth, we parameterizes the quality measure of an ancestry scheme not only by the number of nodes in the given tree but also by its depth. Our main result is the construction of an ancestry scheme that labels $n$-node trees of depth $d$ with labels of size $\log n + 2 \log d + O(1)$. This result is essentially optimal for trees of constant depth which is in fact the typical case in real XML data. In addition to our main result, we prove a result that may be of independent interest concerning the existence of a *universal graph* of linear size for the family of constant depth trees. This result improves the best known bound of $2^{O(\log^* n)} n$ (Alstrup and Rauhe, FOCS 02) on the size of such a universal graph (the latter bound, however, holds for a universal graph for the family of all $n$-node trees).

5.4.1.1.2. On randomized representations of graphs using short labels

In [13], we consider *Informative labeling schemes* which consist in labeling the nodes of graphs so that queries regarding any two nodes (e.g., are the two nodes adjacent?) can be answered by inspecting merely the labels of the corresponding nodes. Typically, the main goal of such schemes is to minimize the *label size*, that is, the maximum number of bits stored in a label. This concept was introduced by Kannan et al. [STOC'88] and was illustrated by giving very simple and elegant labeling schemes, for supporting adjacency and ancestry queries in $n$-node trees; both these schemes have label size $2 \log n$. Motivated by relations between such schemes and other important notions such as *universal* graphs, extensive research has been made by the community

to further reduce the label sizes of such schemes as much as possible. The current state of the art *adjacency* labeling scheme for trees has label size $\log n + O(\log^* n)$ by Alstrup and Rauhe [FOCS'02], and the best known *ancestry* scheme for (rooted) trees has label size $\log n + O(\sqrt{\log n})$ by Abiteboul et al., [SICOMP 2006]. We investigate the above notions from a probabilistic point of view. Informally, the goal is to investigate whether the label sizes can be improved if one allows for some probability of mistake when answering a query, and, if so, by how much. For that, we first present a model for probabilistic labeling schemes, and then construct various probabilistic one-sided error schemes for the adjacency and ancestry problems on trees. Some of our schemes significantly improve the bound on the label size of the corresponding deterministic schemes, while the others are matched with appropriate lower bounds showing that, for the resulting guarantees of success, one cannot expect to do much better in term of label size.

*5.4.1.2. Compact labeling for spanning trees*

**Participants:** Reuven Cohen [Weizmann Institute of Science, Israel], Pierre Fraigniaud, David Ilcinkas [CNRS LABRI, University of Bordeaux, France], Amos Korman, David Peleg [Weizmann Institute of Science, Israel].

The article [17] (awarded the best conference paper) deals with compact label-based representations for trees. Consider an $n$-node undirected connected graph $G$ with a predefined numbering on the ports of each node. The all-ports tree labeling $\mathrm{cL}_{all}$ gives each node $v$ of $G$ a label containing the port numbers of all the tree edges incident to $v$. The upward tree labeling $\mathrm{cL}_{up}$ labels each node $v$ by the number of the port leading from $v$ to its parent in the tree. Our measure of interest is the worst case and total length of the labels used by the scheme, denoted $M_{up}(T)$ and $S_{up}(T)$ for $\mathrm{cL}_{up}$ and $M_{all}(T)$ and $S_{all}(T)$ for $\mathrm{cL}_{all}$. The problem studied in this paper is the following: Given a graph $G$ and a predefined port labeling for it, with the ports of each node $v$ numbered by $0, ..., \deg(v) - 1$, select a rooted spanning tree for $G$ minimizing (one of) these measures. We show that the problem is polynomial for $M_{up}(T)$, $S_{up}(T)$ and $S_{all}(T)$ but NP-hard for $M_{all}(T)$ (even for 3-regular planar graphs). We show that for every graph $G$ and port numbering there exists a spanning tree $T$ for which $S_{up}(T) = O(n \log\log n)$. We give a tight bound of $O(n)$ in the cases of complete graphs with arbitrary labeling and arbitrary graphs with symmetric port assignments. We conclude by discussing some applications for our tree representation schemes.

*5.4.1.3. Labeling schemes for trees with queries*

**Participants:** Amos Korman, Shay Kutten [Technion, Israel].

On a more prospective way, [18] is intended more to ask questions than to give answers. Specifically, we consider models for labeling schemes, and discuss issues regarding the number of labels consulted vs. the sizes of the labels.

Recently, quite a few papers studied methods for representing network properties by assigning *informative labels* to the vertices of a network. Consider some graph function $f$ on pairs of vertices (for example, $f$ can be the distance function). In an $f$-labeling scheme, the labels are constructed in such a way so that given the labels of any two vertices $u$ and $v$, one can compute the function $f(u, v)$ (e.g. the graph distance between $u$ and $v$) just by looking at these two labels. Some very involved lower bounds for the sizes of the labels were proven. Also, some highly sophisticated labeling schemes were developed to ensure short labels.

In this paper, we demonstrate that such lower bounds are very sensitive to the number of vertices consulted. That is, we show several constructions of such labeling schemes that beat the lower bounds by large margins. Moreover, as opposed to the strong technical skills that were needed to develop the traditional labeling schemes, most of our schemes are almost trivial. The catch is that in our model, one needs to consult the labels of *three* vertices instead of two. That is, a query about vertices $u$ and $v$ can access also the label of some third vertex $w$ ($w$ is determined by the labels of $u$ and $v$). More generally, we address the model in which a query about vertices $u$ and $v$ can access also the labels of $c$ other vertices. We term our generalized model *labeling schemes with queries*.

The main importance of this model is theoretical. Specifically, this paper may serve as a first step towards investigating different tradeoffs between the amount of labels consulted and the amount of information stored at each vertex. As we show, if all vertices can be consulted then the problem almost reduces to the corresponding sequential problem. On the other hand, consulting just the labels of $u$ and $v$ (or even just the label of $u$) reduces the problem to a purely distributed one. Therefore, in a sense, our model spans a range of intermediate notions between the sequential and the distributed settings.

In addition to the theoretical interest, we also show cases that schemes constructed for our model can be translated to the traditional model or to the sequential model, thus, simplifying the construction for those models as well. For implementing query labeling schemes in a distributed environment directly, we point at a potential usage for some new paradigms that became common recently, such as P2P and overlay networks.

### 5.4.2. *Efficient graph spanners*

#### *5.4.2.1. Distributed Computation of Sparse Spanner*

**Participants:** Bilel Derbel [CNRS LIFL, University of Lille, France], Cyril Gavoille [CNRS LABRI, University of Bordeaux, France], David Peleg [Weizmann Institute of Science, Israel], Laurent Viennot.

An $(\alpha, \beta)$-spanner of a graph $G$ is a subgraph $H$ that approximates distances in $G$ within a multiplicative factor $\alpha$ and an additive error $\beta$. More precisely, for any two nodes $u, v$ of $G$, $d_H(u, v) \leq \alpha \cdot d_G(u, v) + \beta$. Computing sparse spanners is a fundamental problem of distributed computing and compact routing.

In [9] we provide distributed algorithms for computing sparse $(\alpha, \beta)$-spanners with $\alpha$ close to one. We present a generic distributed algorithm that in constant number of rounds computes, for every $n$-node graph and integer $k \geq 1$, an $(\alpha, \beta)$-spanner of $O(\beta n^{1+1/k})$ edges, where $\alpha$ and $\beta$ are constants depending on $k$. For suitable parameters, this algorithm provides a $(2k-1, 0)$-spanner of at most $kn^{1+1/k}$ edges in $k$ rounds. For $k = 2$ and constant $\varepsilon > 0$, it can also produce a $(1+\varepsilon, 2-\varepsilon)$-spanner of $O(n^{3/2})$ edges in constant time. More interestingly, for every integer $k > 1$, it can construct in constant time a $(1+\varepsilon, O(1/\varepsilon)^{k-2})$-spanner of $O(\varepsilon^{-k+1} n^{1+1/k})$ edges. Such deterministic construction was not previously known.

We also present a second generic deterministic and distributed algorithm based on the construction of small dominating sets and maximal independent sets. After computing such sets in sub-polynomial time, it constructs at its best a $(1+\varepsilon, \beta)$-spanner with $O(\beta n^{1+1/k})$ edges, where $\beta = k^{\log(\log k/\varepsilon)+O(1)}$. For $k = 3$, it provides a $(1+\varepsilon, 6-\varepsilon)$-spanner with $O(\varepsilon^{-1} n^{4/3})$ edges.

#### *5.4.2.2. Spanners for Ad Hoc Networks*

**Participants:** Philippe Jacquet [INRIA Hipercom, CRI Paris Rocquencourt, France], Laurent Viennot.

In [16], we study a variation of spanners where a node is supposed to know the list of its neighbors. This is particularly suited to optimize the set of links advertized in a practical link state algorithm running in a possibly dense network such as an ad hoc network. Given an unweighted graph $G$, a sub-graph $H$ with vertex set $V(H) = V(G)$ is an $(a, b)$-*remote-spanner* if for each pair of points $u$ and $v$ the distance between $u$ and $v$ in $H_u$, the graph $H$ augmented by all the edges between $u$ and its neighbors in $G$, is at most $a$ times the distance between $u$ and $v$ in $G$ plus $b$. We extend this definition to $k$-connected graphs by considering the minimum length sum over $k$ disjoint paths as a distance. We then say that an $(a, b)$-remote-spanner is $k$-*connecting*.

We give distributed algorithms for computing $(1+\varepsilon, 1-2\varepsilon)$-remote-spanners for any $\varepsilon > 0$, $k$-connecting $(1, 0)$-remote-spanners for any $k \geq 1$ (yielding $(1, 0)$-remote-spanners for $k = 1$) and 2-connecting $(2, -1)$-remote-spanners. All these algorithms run in constant time for any unweighted input graph. The number of edges obtained for $k$-connecting $(1, 0)$-remote-spanner is within a logarithmic factor from optimal (compared to the best $k$-connecting $(1, 0)$-remote-spanner of the input graph). Interestingly, sparse $(1, 0)$-remote-spanners (i.e. preserving exact distances) with $O(n^{4/3})$ edges exist in random unit disk graphs. The number of edges obtained for $(1+\varepsilon, 1-2\varepsilon)$-remote-spanners and 2-connecting $(2, -1)$-remote-spanners is linear if the input graph is the unit ball graph of a doubling metric (even if distances between nodes are unknown). Our methodology consists in characterizing remote-spanners as sub-graphs containing the union of small depth tree sub-graphs dominating nearby nodes. This leads to simple local distributed algorithms.

In [15], we analyse the size of $k$-connecting (1,0)-remote- spanners in classical random graph models. Interestingly, the expected compression ratio in number of edges is $O(\frac{k}{n}\log n)$ in the Erdös-Rényi graph model and $O((\frac{k}{n})^{\frac{2}{3}})$ in the unit disk graph model with a uniform Poisson distribution of nodes.

This work gives a theoretical foundation to the OLSR routing protocol (RFC 3626). In particular, it shows that multipoint relays (which are the basis of OLSR functionning) are an inherent structure for providing $(1,0)$-remote-spanners, i.e. optimal routes.

### *5.4.3. Graph decompositions*

**Participants:** Binh-Minh Bui-Xuan [University of Bergen, Norway], Michel Habib, Vincent Limouzy, Fabien de Montgolfier, Michael Rao [CNRS LABRI, University of Bordeaux, France].

#### *5.4.3.1. NLC decomposition*

Many *width* graph decompositions have been proposed. Thanks to Courcelle theorem, they allow to efficiently solve many hard (NP-complete) problems for graph classes, provided the decomposition width is bounded. NLC decomposition is a variation of cliquewidth, where the decomposition is a labelled tree. In [48], the recognition of graphs of NLC 2 is addressed. The previous time complexity is improved to $O(n^2m)$, and the algorithm is robust.

#### *5.4.3.2. Umodular decomposition*

A new decomposition of combinatorial structures is presented in [33]. Is is based on a generalisation of the modular decomposition. When applied to undirected graph, it gives the bijoin decomposition, and when applied to tournaments, it gives a new decomposition. We present proofs of existence and uniqueness of a decomposition tree, and polynomial-time algorithms.

## 5.5. Discrete Optimization Algorithms

### *5.5.1. Efficient encoding of pseudo-boolean constraints*

**Participants:** Olivier Bailleux [LERSIA, University of Bourgogne], Yacine Boufkhad, Olivier Roussel [CNRS CRIL, University of Artois].

In [4] the open question of the existence of a polynomial size CNF encoding of pseudo-Boolean (PB) constraints such that generalized arc consistency (GAC) is maintained through unit propagation (UP) is answered affirmatively. All previous encodings of PB constraints either didn't allow UP to maintain GAC, or were of exponential size in the worst case. In the above cited paper an encoding that realizes both of the desired properties is presented. From a theoretical point of view, this narrows the gap between the expressive power of clauses and the one of pseudo-Boolean constraints.

### *5.5.2. Three way decomposition of permutation problems*

**Participants:** Dominique Fortin, Ider Tseveendorj.

Given a $m \times m$ flow matrix $F$ and a $n \times n$ distance matrix $D$, the Quadratic Assignment Problem (QAP) aims at minimizing the overall energy to carry the flows among the facilities assigned to locations related by the distance matrix; using a binary assignment $m \times n$ matrix $X$, its formulation is minimizing:

$$QAP(D,F) = \langle F, XDX^t \rangle \text{s.t. } X \in \mathfrak{P} \tag{1}$$

where $\mathfrak{P}$ denotes, loosely speaking, the set of permutations. In more standard notations, when $m < n$

$$\begin{cases} \min & \langle F, XDX^t \rangle \\ \text{s.t.} & \sum_l X_{fl} = 1 \\ & \sum_f X_{fl} \leq 1 \end{cases} \tag{2}$$

Many practical problems give rise to $(QAP)$; among special cases, the Traveling Salesperson Problem (TSP) corresponds to $F = I$ the identity matrix. In order to guess the correlation structure between constraints and objective in 0-1 programming, we devised in [2] a method to firstly sample the distribution of fractional solutions of the continuous relaxation and then use this distribution to select an effective branching rule to early detect a *good solution*; in this sense, it will prune many nodes in the branching tree. Computational results on multiknapsack and the maximum clique problem prove efficiency of this adaptative approach independent of the given problem. So, it was tempting to experiment this approach to more general integer programming especially those dealing with permutation such as $(QAP)$ [1], [25], [24].

# 6. Contracts and Grants with Industry

## 6.1. Collaboration with France Telecom (CRC Mardi)

**Participants:** Diego Perino, Dmitry Lebedev [France Telecom R&D, Issy Les Moulineaux, France], Fabien Mathieu [France Telecom R&D, Issy Les Moulineaux, France], Fabien de Montgolfier, Julien Reynier, Laurent Viennot, Simon Gwendal [France Telecom R&D, Issy Les Moulineaux, France].

MARDI is a collaboration contract between Inria and France Telecom. It gathers Gang and Spontex (FT) around the study of decentralized networks over Internet. Spontex is a transversal project on cooperative networks. Diego Perino is funded through this collaboration and co-supervised by Fabien Mathieu and Laurent Viennot.

### *6.1.1. Measuring Internet with and for peer-to-peer networks*

A first aspect of the project consist in studying Internet latencies in order to understand how logical overlays can be optimized with respect to delays. A possible track for gathering valuable large scale measures is to use a peer-to-peer network for measuring latencies. Interestingly, it is possible to find shortcuts in the Internet where the route through a relay can be faster than the direct route.

### *6.1.2. Modeling node and connection dynamics*

This item is connected to the affinity model where peers tend to connect preferentially to some peers based on some measured or inferred criteria. Connecting peers according to delays is a special case of affinity where a peer connects preferentially to peers with low RTT. Additional properties can be proven for this case to prove the convergence of a dynamic system following this low RTT strategy.

### *6.1.3. Peer-to-peer application design*

The third part of the project aims at designing efficient structuring algorithm for decentralized applications. It relies on the previous parts. Measuring and modeling Internet latencies can be used to obtain a first coarse solution to a fast overlay, and the affinity models can be use to tune the solution and to adapt it under node churn.

# 7. Other Grants and Activities

## 7.1. National initiatives

### *7.1.1. ANR Algorithm Design and Analysis for Implicitly and Incompletely Defined Interaction Networks (ALADIN)*

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

Pierre Fraigniaud is leading an ANR project “blanc” (i.e. fundamental research) about the fundamental aspects of large interaction networks enabling massive distributed storage, efficient decentralized information retrieval, quick inter-user exchanges, and/or rapid information dissemination. The project is mostly oriented towards the design and analysis of algorithms for these (logical) networks, by taking into account proper ties inherent to the underlying infrastructures upon which they are built. The infrastructures and/or overlays considered in this project are selected from different contexts, including communication networks (from Internet to sensor networks), and societal networks (from the Web to P2P networks).

## 7.2. Actions Funded by the EC

### *7.2.1. COST 295 – Dynamo (2005-2009)*

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.

# 8. Dissemination

## 8.1. Services to the scientific community

- 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 [28].
- 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).
- Pierre Fraigniaud has participated to the program committees of:

  PODC 2009 28th ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, Calgary, Canada, August 10-12, 2009.

  MFCS 2009 34th International Symposium on Mathematical Foundations of Computer Science, August 24 - 28, 2009 Novy Smokovec, Slovakia.

  ICALP 2009 36th International Colloquium on Automata, Languages and Programming, Track C "Foundations of Networked Computation", July 5 - 12, 2009 Rhodes, Greece.

  ESA 2009 17th Annual European Symposium on Algorithms, Sept. 7-9, 2009, Copenhagen.

  IPDPS 2009 23rd IEEE International Parallel and Distributed Processing Symposium, Rome, Italy, May 25-29, 2009.

  ICDCN 2009 10th International Conference on Distributed Computing and Networking, January 3-6, 2009, Gachibowli, Hyderabad, India.

  SSS 2009 11th International Symposium on Stabilization, Safety, and Security of Distributed Systems, Lyon, France, November 3-6, 2009.

## 8.2. Teaching

Master MPRI Michel Habib is in charge of a course entitled “graph algorithms”. Pierre Fraigniaud is in charge of the course “Algorithmique distribuée pour les réseaux”.

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

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

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

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

## 8.3. Theses

### *8.3.1. Defended theses*

- Diego Perino *On Resource allocation algorithms for peer-to-peer multimedia streaming.* (CIFRE Orange Labs). Defended November 16th 2009.

  Jury: Advisors (Directeurs): Fabien Mathieu (Orange Labs), Laurent Viennot (INRIA);

  Reviewers (Rapporteurs): Marco Ajmone Marsan (Politecnico di Torino), Pascal Felber (Université de Neuchatel), Anne-Marie Kermarrec (INRIA);

  Examiners (Examinateurs): Pierre Fraignaud (CNRS), Arnaud Legout (INRIA), Laurent Massoulié, (Thomson Lab).

### *8.3.2. Ongoing theses*

- Damien Noguës: *Algorithmes pour les graphes $\delta$-hyperboliques* (AMN)
- Mauricio Soto *Algorithmes de pair à pair et analyse de la structure d'Internet* (Chile-France Allocation).
- Anh Hoang Phan *Overlays structurés en pair à pair* (BDI).
- Hien Hu To *Décomposition de graphes* (AMX).

# 9. Bibliography

## Year Publications

### Articles in International Peer-Reviewed Journal

[1] D. FORTIN, I. TSEVEENDORJ. *A Trust Branching Path Heuristic for permutation Problems*, in "International Journal of Pure and Applied Mathematics", vol. 56, n$^{o}$ 3, 2009, p. 329-343.

[2] D. FORTIN, I. TSEVEENDORJ. *A Trust Branching Path Heuristic for Zero-One Programming*, in "European Journal of Operational Research", vol. 197, n$^{o}$ 2, 2009, p. 439-445.

[3] M. KIWI, M. SOTO, C. THRAVES. *Adversarial queuing theory with setups*, in "Theor. Comput. Sci.", vol. 410, n$^{o}$ 8-10, 2009, p. 670-687, http://dx.doi.org/10.1016/j.tcs.2008.09.064CL.

### International Peer-Reviewed Conference/Proceedings

[4] O. BAILLEUX, Y. BOUFKHAD, O. ROUSSEL. *New Encodings of Pseudo-Boolean Constraints into CNF*, in "Theory and Applications of Satisfiability Testing - SAT 2009, Royaume-Uni Swansea", 2009, p. 181-194, http://hal.archives-ouvertes.fr/hal-00415126/en/.

[5] H. BAUMANN, P. CRESCENZI, P. FRAIGNIAUD. *Parsimonious Flooding in Dynamic Graphs*, in "Proc. 28th ACM Symposium on Principles of Distributed Computing (PODC)", 2009, p. 260-269.

[6] H. BAUMANN, P. FRAIGNIAUD. *Sub-linear Universal Spatial Gossip Protocols*, in "Proc. 16th International Colloquium on Structural Information and Communication Complexity (SIROCCO)", LNCS, Springer, 2009.

[7] Y. BOUFKHAD, F. MATHIEU, F. DE MONTGOLFIER, D. PERINO, L. VIENNOT. *An Upload Bandwidth Threshold for Peer-to-Peer Video-on-Demand Scalability*, in "23rd IEEE International Parallel and Distributed Processing Symposium (IPDPS)", 2009, p. 1–10.

[8] Y. BOUFKHAD, F. MATHIEU, F. DE MONTGOLFIER, D. PERINO, L. VIENNOT. *Fine Tuning of a Distributed VoD System*, in "18th IEEE International Conference on Computer Communications and Networks (ICCCN)", 2009, p. 1–7.

[9] B. DERBEL, C. GAVOILLE, D. PELEG, L. VIENNOT. *Local Computation of Nearly Additive Spanners*, in "DISC, Espagne Elche/Elx", vol. 5805, Springer, 2009, p. 176-190, http://hal.archives-ouvertes.fr/hal-00408725/en/IL.

[10] Y. EMEK, P. FRAIGNIAUD, A. KORMAN, 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 IL .

[11] Y. EMEK, P. FRAIGNIAUD, A. KORMAN, A. ROSEN. *Online Computation with Advice*, in "Proc. 36th International Colloquium on Automata, Languages and Programming (ICALP)", LNCS 5555, 2009, p. 427-438 IL .

[12] P. FRAIGNIAUD, G. GIAKKOUPIS. *The Effect of Power-Law Degrees on the Navigability of Small Worlds*, in "Proc. 28th ACM Symposium on Principles of Distributed Computing (PODC)", 2009, p. 240-249.

[13] P. FRAIGNIAUD, A. KORMAN. *On Randomized Representations of Graphs Using Short Labels*, in "Proc. 21st ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)", 2009, p. 131-137 IL .

[14] P. FRAIGNIAUD, A. KORMAN. *Compact Ancestry Labeling Schemes for XML Trees*, in "21st ACM-SIAM Symposium on Discrete Algorithms (SODA)", 2010, to appear IL .

[15] P. JACQUET, L. VIENNOT. *Average Size of Unstretched Remote-Spanners*, in "5th ACM Workshop on Analytic Algorithmics and Combinatorics (ANALCO)", 2009, p. 23–33.

[16] P. JACQUET, L. VIENNOT. *Remote spanners: what to know beyond neighbors*, in "23rd IEEE International Parallel and Distributed Processing Symposium (IPDPS)", 2009, p. 1–10.

[17] A. KORMAN. *Compact Routing Schemes for Dynamic Trees in the Fixed Port Model*, in "10th International Conference on Distributed Computing and Networking, (ICDCN)", Lecture Notes in Computer Science, vol. 5408, 2009, p. 218-229.

[18] A. KORMAN, S. KUTTEN. *A Note on Models for Graph Representations*, in "Theoretical Computer Science(TCS)", vol. 410, 2009, p. 1401-1412 IL .

[19] L. MUSCARIELLO, D. 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.

[20] L. MUSCARIELLO, D. PERINO. *Modeling multi-path routing and congestion control under FIFO and Fair Queuing*, in "IEEE Conference on Local Computer Network (LCN)", 2009.

[21] L. MUSCARIELLO, D. PERINO, B. NARDELLI. *Towards real implementations of dynamic robust routing paradigms exploiting path diversity*, in "IEEE Conference on Ultra Modern Telecommunications (ICUMT)", 2009.

[22] L. MUSCARIELLO, D. PERINO, D. ROSSI. *Do Next Generation Networks need Path Diversity?*, in "International Conference on Communications (ICC)", 2009.

### National Peer-Reviewed Conference/Proceedings

[23] N. HEGDE, F. MATHIEU, D. PERINO. *Diffusion epidemique de chunks en quasi-direct : la taille compte*, in "Algotel", 2009.

### Scientific Books (or Scientific Book chapters)

[24] D. FORTIN. *Hadamard's matrices, Grothendieck's constant and square Root Two*, in "Optimal Control: Theory and Applications", Springer, 2009.

[25] D. FORTIN, I. TSEVEENDORJ. *Escaping from a local solution region of convex maximization problem*, in "Global Optimization : Theory, Methods & Applications I", lecture notes in decision sciences, vol. 12(B), Global-Link, september 2009, p. 943-948.

### Research Reports

[26] N. HEGDE, F. MATHIEU, D. PERINO. *Size does matter (in P2P live streaming)*, INRIA, 2009, RR-7032, Rapport de recherche.

[27] F. MATHIEU, D. PERINO. *On Resource Aware Algorithms in Epidemic Live Streaming*, INRIA, 2009, http://hal.inria.fr/inria-00414706/en/, RR-7031, Rapport de recherche.

### Scientific Popularization

[28] L. VIENNOT. *Idée reçue : Web et Internet, c'est la même chose*, 2009, http://interstices.info/jcms/c_42755/web-et-internet-c-est-la-meme-chose, Interstices (French vulgarization site).

## References in notes

[29] I. ABRAHAM, C. GAVOILLE, A. V. GOLDBERG, D. MALKHI. *Routing in Networks with Low Doubling Dimension*, in "26th International Conference on Distributed Computing Systems (ICDCS)", IEEE Computer Society Press, July 2006.

[30] L. ADAMIC, R. LUKOSE, A. PUNIYANI, B. HUBERMAN. *Search in power-law networks*, in "Physical Review E", 2001.

[31] R. ALBERT, H. JEONG, A.-L. BARABASI. *Diameter of the World Wide Web*, in "Nature", vol. 401, 1999, p. 130–131.

[32] P. ASSOUAD. *Plongements lipschitziens dans Rn*, in "Bull. Soc. Math. France", 1983.

[33] B.-M. BUI XUAN, M. HABIB, V. LIMOUZY, F. DE MONTGOLFIER. *Unifying two graph decompositions with modular decomposition*, in "The 18th International Symposium on Algorithms and Computation (ISAAC 2007)", 2007.

[34] B. COHEN. *Incentives Build Robustness in BitTorrent*, in "Workshop on Economics of Peer-to-Peer Systems", 2003.

[35] W. COOK, P. SEYMOUR. *Tour Merging via Branch-Decomposition*, in "INFORMS Journal on Computing", 2003.

[36] B. COURCELLE, J. MAKOWSKY, U. ROTICS. *Linear time solvable optimization problems on graphs of bounded clique-width*, in "Theory Comput. Syst.", vol. 33(2), 2000, p. 125–150.

[37] J. DAI, O. JENNINGS. *Stabilizing Queueing Networks with Setups*, in "Mathematics of Operation Research", vol. 29, 2004.

[38] P. DUCHON, N. HANUSSE, E. LEBHAR, N. SCHABANEL. *Could any graph be turned into a small-world?*, in "DISC: International Symposium on Distributed Computing", vol. 355(1), 2006, p. 96-103.

[39] A. FABRIKANT, E. KOUTSOUPIAS, C. PAPADIMITRIOU. *Heuristically optimized trade-offs: A new paradigm for power laws in the Internet*, in "ICALP", 2002.

[40] P. FRAIGNIAUD. *A new perspective on the small-world phenomenon: Greedy routing in tree-decomposed graphs*, in "Proceedings of the 13th Annual European Symposium on Algorithms (ESA)", 2005, p. 791–802.

[41] P. FRAIGNIAUD, E. LEBHAR, Z. LOTKER. *A Doubling Dimension Threshold* $\Theta(\log\log n)$ *for Augmented Graph Navigability*, in "Proceedings of14th Annual European Symposium (ESA)", 2006, p. 376-386.

[42] ANH-TUAN. GAI, L. VIENNOT. *Broose: A Practical Distributed Hashtable based on the De-Brujin Topology*, in "The Fourth IEEE International Conference on Peer-to-Peer Computing", 2004.

[43] ANH-TUAN. GAI, L. VIENNOT. *Incentive, Resilience and Load Balancing in Multicasting through Clustered de Bruijn Overlay Network (PrefixStream)*, in "14th IEEE International Conference on Networks", September 2006.

[44] D. GAMARNIK. *Stability of adversarial queues via fluid model*, in "Proc. of the 39th Annual Symposium on Foundations of Computer Science, Palo Alto, CA", 1998, p. 60–70.

[45] D. KARGER, M. RUHL. *Finding nearest neighbors in growth-restricted metrics*, 2002.

[46] J. KLEINBERG. *The Small-World Phenomenon: An Algorithmic Perspective*, in "Proceedings of the 32nd ACM Symposium on Theory of Computing (STOC)", 2000, p. 163–170.

[47] R. KRAUTHGAMER, J. LEE. *Navigating nets: Simple algorithms for proximity search*, 2004.

[48] V. LIMOUZY, F. DE MONTGOLFIER, M. RAO. *NLC-2 graph recognition and isomorphism*, in "WG'07, 33rd International Workshop on Graph-Theoretic Concepts in Computer Science", 2007.

[49] P. MAYMOUNKOV, D. MAZIERES. *Kademlia: A peer-to-peer information system based on the xor metric*, in "1st International Workshop on Peer-to-Peer Systems (IPTPS)", 2002.

[50] S. RATNASAMY, P. FRANCIS, M. HANDLEY, R. KARP, S. SHENKER. *A scalable content-addressable network*, in "Proc. ACM SIGCOMM'01, San Diego, CA", August 2001.

[51] N. ROBERSTON, P. SEYMOUR. *Graph Minors X. Obstructions to tree-decomposition*, in "J. Combinatorial Theory Ser. B", vol. 52(2), 1991, p. 153–190.

[52] A. ROWSTRON, Y. DRUSCHEL. *Pastry: Scalable, Decentralized Object Location, and Routing for Large-Scale Peer-to-Peer Systems*, in "IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)", LNCS 2218, 2001, p. 329–350.

[53] A. SLIVKINS. *Distance Estimation and Object Location via Rings of Neighbors*, in "Proceedings of the 24th Annual ACM Symposium on Principles of Distributed Computing (PODC)", 2005, p. 41-50.

[54] I. STOICA, R. MORRIS, D. KARGER, M. KAASHOEK, H. BALAKRISHNAN. *Chord: A scalable peer-to-peer lookup service for internet applications*, in "Proceedings of the 2001 conference on applications, technologies, architectures, and protocols for computer communications (SIGCOMM)", ACM Press, 2001, p. 149–160.

[55] D. WATTS, S. STROGATZ. *Collective dynamics of small-world networks*, 1998.