Coordinator's names	Carole DELPORTE-GALLET Hugues FAUCONNIER *LIAFA-GANG University Paris Diderot*
Acronym	DISPLEXITY
Titre de la proposition du projet	Calculabilité et complexité en distribué
Proposal title	Distributed Computing : computability and complexity
Evaluation committee	SIMI 2- Science informatique et applications
Multidisciplinary proposal	NO
Type of research	Basic Research
International cooperation	NO
Grant requested **Proposal duration**	829 515, 47 4 years

Contents

1 Proposal Abstract **3**

2 Context, Positionning and Objectives of the proposal **4**
2.1 Context of the proposal . . . 4
2.2 Objectives, originality and/or novelty of the proposal . . . 6

3 Scientific and technical programme, proposal organisation **6**
3.1 Scientific programme, proposal structure . . . 6
3.2 Description by task . . . 8
3.2.1 Task 1: Project Management . . . 8
3.2.2 Task 2: "yes-no" and Decision Problems in Distributed Computing . . . 8
3.2.3 Task 3: Oracles . . . 9
3.2.4 Task 4: Complexity classes . . . 11
3.2.5 Task 5: Non-determinism in distributed computing . . . 12
3.2.6 Task 6: New computational paradigms/frameworks . . . 13
3.3 Tasks schedule, deliverables and milestones . . . 14

4 Dissemination and exploitation of results, intellectuel property **15**

5 Consortium description **16**
5.1 Partners description and relevance, complementarity . . . 16
5.2 Qualification of the proposal coordinator . . . 17
5.3 Qualification and contribution of each partner . . . 18

6 Scientific justification of requested ressources **21**
6.1 Partner 1: Paris . . . 21
6.2 Partner 2: Rennes . . . 22
6.3 Partner 3: Bordeaux . . . 23

7 Annexes **24**
7.1 References . . . 24
7.2 CV, resume . . . 26
7.3 Staff involvment in other contracts . . . 44

1 Proposal Abstract

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

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

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

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

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

- Formalizing yes/no-problems (decision problems) in the context of distributed computing. Such problems are expected to play an analogous role in the field of distributed computing as that played by decision problems in the context of sequential computing.

- Formalizing decision problems (yes/no-problems) in the context of distributed computing. Such problems are expected to play an analogous role in the field of distributed computing as that played by decision problems in the context of sequential computing.

- Revisiting the various explicit (e.g., failure-detectors) or implicit (e.g., a priori information) notions of oracles used in the context of distributed computing allowing us to express them in terms of decidability/complexity classes based on oracles.

- Identifying the impact of non-determinism on complexity in distributed computing. In particular, DISPLEXITY aims at a better understanding of the apparent lack of impact of non-determinism in the context of fault-tolerant computing, to be contrasted with the apparent huge impact of non-determinism in the context of network computing. Also, it is foreseen that non-determinism will enable the comparison of complexity classes defined in the context of fault-tolerance with complexity classes defined in the context of network computing.

- Last but not least, DISPLEXITY will focus on new computational paradigms and frameworks, including, but not limited to distributed quantum computing and algorithmic game theory (e.g., network formation games).

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

2 Context, Positionning and Objectives of the proposal

2.1 Context of the proposal

The distributed community can be viewed as the union of two sub-communities. Even though they are not completely disjoint, they are disjoint enough not to leverage each other's results. At a high level, one is mostly interested in timing issues (clock drifts, link delays, crashes, etc.) while the other one is mostly interested in spatial issues (network structure, memory requirements, etc.). Indeed, one sub-community is mostly focusing on the combined impact of asynchronism and faults on distributed computation, while the other addresses the impact of network structural properties on distributed computation. Both communities address various forms of computational complexities, through the analysis of different concepts. This includes, e.g., failure detectors and wait-free hierarchy for the former community, and compact labeling schemes and computing with advice for the latter community.

We illustrate these points by some examples below.

In the wait-free model, each process starts with a private input value, and decides irrevocably a private output value after a finite number of steps. The processes communicate with each other through shared objects, but the output decision of each process must occur independently of the speed or actions of other processes. The output value decided by each process, which must satisfy some input-output specification, could however depend on the actions of the other processes. A wait-free distributed algorithm performs correctly in any asynchronous environment, and, in particular, it tolerates an arbitrary number of process crashes. In this way, the wait-free model may be considered as a special case of t-resiliency for t equal to the number of processes. Some simulations, like in [4, 5] and general results prove that t-resiliency may generally be reduced to the wait-free model.

Not all tasks (specified by input-output relations) are wait-free solvable. More specifically, the famous FLP Theorem [16] states that consensus is not 1-resilient and then not wait-free solvable. Several research frameworks in the context of wait-free and more generally fault-tolerant computing offer similarities with computing with oracle machines. One typical example is provided by the failure detector theory [8] which is based on oracles providing nodes with limited information about the failures. Failure detectors may be compared by reduction [7]. Hence if every failure detector enabling to solve task A may be reduced to a failure detector enabling to solve task B, then task A is harder (i.e. needs mode information about failure) than task B. In this way, failure detectors enable to define hierarchy among unsolvable tasks [10].

Another example is provided by the wait-free hierarchy (including Herlihy's hierarchy [22]) which maps object to positive levels such that an object is at level n if and only if it offers some form of universality for a system of n processes in the wait-free model.

In network computing, networks are modeled by graphs $G = (V; E)$, in which messages are exchanged between the processes (i.e., the nodes in V), along the links of the network (i.e., the edges in E). The model $\mathcal{LOCAL}$ [27] assumes that the nodes operate in synchronous rounds, where a round enables to exchange a message of arbitrarily large size between every pair of nodes linked by an edge. Hence, in the $\mathcal{LOCAL}$ model, the main measure of interest is the maximum distance at which processes interact. A task (e.g., coloring the nodes of properly with at most $\Delta+1$ colors) is local if and only if it has an algorithm solving it in $O(1)$ rounds in the $\mathcal{LOCAL}$ model[1]. Not all problems are local. In particular, Linial [25] proved that graph coloring is not local, even in rings ($\mathcal{O}(\log n)$ rounds are required to 3-color the ring). Naor and Stockmeyer [26] studied the class LCL of locally checkable labeling problems. As in the case of wait-free computation, several research frameworks in the context of local computing offer similarities with computing with oracle machines. One typical example is distributed computing with advice, which is aiming at capturing the impact of global structural knowledge (e.g., the size of the network) on the efficiency of solving distributing computing tasks (like coloring or broadcasting). Another example is proof labeling schemes [24], which is aiming at providing each node with labels such that the consistency of distributed data-structures (e.g., MST, spanning tree, etc.) can be checked in one round, just by exchanging labels between neighbors.

[1] An interesting restriction do the $\mathcal{LOCAL}$ model is the $\mathcal{CONGEST}(B)$ model in which the size of exchanged messages is bounded by B.

Although the examples mentioned above offer the flavor of computational complexity, they are not quite clearly positioned as part of computational complexity. Comparing problems (or, more appropriately, tasks) that are not solvable in the fault-tolerance models is achieved through the analysis of their weakest failure detectors, while wait-free reductions are rather performed between objects and not between problems. Similarly, the concepts of advices and proof labeling schemes enable to handle and sometime compare problems that are not locally computable. Yet the papers addressing these notions are mostly problem oriented, and do not provide a global consistent distributed computing complexity framework. It is interesting to note that although the above examples are dealing with computational complexity, each time, the approach is ad hoc. We believe that we reach the point where it is necessary to base a general theory of computability and complexity for distributed computing. Such a theory is crucial to have a better understanding of what is distributed computing. Considering the relationship between classical algorithmic and classical complexity theory, it is clear that it can be expected that computability and complexity theory for distributed computing will have many applications concerning distributed algorithms. The main goal of the project is precisely to work towards this goal.

This project is clearly very ambitious. Yet, more than 30 years of distributed computing have enabled to accumulate a very strong knowledge and the partners of this project are at the leading edge of research on distributed computing and cover a large spectrum of this discipline. Should this project successful, we foresee a significant impact in the area and a unique position of the French research.

Related ANR project

Some partners (Bordeaux and Paris) participate to the ALADIN ANR Project (2007-2011) that aimed at studying fundamental aspects of large interaction networks. More specifically, one of the topics of ALADIN is to deal with the design of distributed algorithms with limited knowledge and/or limited probing capacity. The ALADIN project will be finished at the end of the year (2011). The works realized by the partners in this ANR will be helpfull for the project. In particular, some objectives of the DISPLEXITY project comes from recent results [19, 20] of ALADIN concerning $\mathcal{LOCAL}$ distributed decision and decidability classes for mobile agents computing.

Some partners (Paris and Rennes) participate to the SHAMAN ANR-VERSO (2008-2012) project that focuses on the algorithmic foundations of resource-constrained autonomous large scale systems. The first objective of this project is the design of realistic models encompassing anonymity, dynamism, and/or malicious behavior. This project should give some help concerning some new aspects of distributed computation especially concerning large scale systems. Many works of Rennes concern oracles (failure detectors)(e.g. [2, 3]). In [11], some members of Paris and Rennes give a model that encompasses anonymity and malicious behavior that would be useful for DISPLEXITY and [12] defines the notion of adversaries that will be tackled in this project in the framework of oracles. On the other hand, the progress on computability and complexity made by the DISPLEXITY project will help the SHAMAN project to better evaluate the theoretical power of these models.

Some partners (Bordeaux and Paris) participate to the PROSE ANR Project (Sept 2009-Aug 2012). The goal is the deployment of social networking applications in a delay tolerant manner using opportunistic social contacts as in a peer to peer network, as well as new advanced content recommendation engines. It is a multi-disciplinary project to design opportunistic contact sharing schemes and to characterize the environmental conditions, the usage constraint, as well as the algorithmic and architecture principles that let them operate. One task concerns the theory of dynamic graph and networking modeling that is interesting for DISPLEXITY. In return, PROSE will benefit with help of DISPLEXITY of a better understanding of this model.

2.2 Objectives, originality and/or novelty of the proposal

We have the ambitious project to achieve the reconciliation between the two communities by focusing on the same class of problems, the yes/no-problems, and establishing the scientific foundations for building up a consistent theory of computability and complexity for distributed computing.

The main question addressed 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 [2], 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 [3] as well as Linial's lower bound [25] 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.

3 Scientific and technical programme, proposal organisation

3.1 Scientific programme, proposal structure

The main purpose of this project is to define the basis for the complexity and the computability of distributed computing.

First of all, the classical notions of computability in term of, for example Turing machines, do not apply directly to distributed computing. Distributed computing deals with infinite computations and some models of computations for distributed protocols (e.g. population protocols [1] or self-stabilizing algorithms [15]) are defined in terms of convergence (only computations such that eventually the output stabilizes are considered) and asynchrony in wait-free and fault-tolerant computing are properties on sets of infinite sequences with some fairness conditions. The FLP impossibility result for the consensus [16] does not depend on the computability power of the processes, and this result is then very different from the classical undecidability results with Turing machines.

Concerning complexity problems, distributed computing raises new problems. One of the main point is here the question of the locality of computations as in the $\mathcal{LOCAL}$ [27] or $\mathcal{CONGEST}$ models making the definition of complexity classes relevant for distributed computing necessary.

Nevertheless, even if the classical computability and complexity theory do not apply directly here, it is clear that most of its methods and tools will give the basis to establish a computational and complexity theory for distributed computing. Therefore the main approach of this project is to try to use classical tools

[2] 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.

[3] FLP result [16] is a fundamental result in fault-tolerant distributed computing that proves the impossibility of consensus when at least one process may crash.

of computability and complexity and to adapt them to the specificity of distributed computing. For, this we have identified some of them that will serve as central themes of each task.

As for Turing machine, the main goal of Task 2 is to define "yes-no" problems [4] and then to tackle the relationship between computing and verifying a problem under various computational models for distributed computing. Note that in the distributed context even the definition of "yes-no" problems is not *a priori* obvious depending on who has to decide and several definitions are possible. For distributed computations, the relationship between computing and verifying is not so obvious. For example with faulty processes it may be easier to achieve consensus between processes than verifying that consensus is achieved. Definition and study of "yes-no" problems will enable to tackle this fundamental issue.

Another fundamental tool concerning computability is the notion of oracles, subject of Task 3. Oracles, like failure-detectors [8] in fault-tolerance computing have been introduced to bypass the impossibility result of the consensus. Indeed with some partial information about crashes problems like consensus become possible to solve. Failure detectors encapsulate this information about crashes. They have proved their efficiency enabling to obtain hierarchies between problems impossible to solve in presence of process crashes. The oracles are not only useful in fault-tolerant computing to establish hierarchy between (impossible) problems but, as for classical computability and complexity theory, they are fundamental basic elements of distributed computing. For example, an oracle that gives the number of nodes enable to distinguish between rings of nodes that otherwise are indistinguishable. More generally, the oracles and the notion of reduction with oracles enables to define hierarchy and complexity classes for distributed computing. Hence general study of oracles and reductions in distributed computing will help to define complexity and computability classes for distributed computing.

Defining complexity classes for distributed computing comes naturally after the study of "yes-no" problems and oracles and is the subject of Task 4. Clearly these definitions depend on the definition of the definitions of "yes-no" problems and oracles. Two main kinds of classes are envisioned: complexity classes based on the notion of time (or here number of communication rounds) and classes based on the notion of possibility-impossibility. For example, for the first ones, Local Decision (LD) class is the class of distributed languages with decision in constant time in the $\mathcal{LOCAL}$ model and for the second ones new classes have to be defined for the wait-free and fault-tolerant models. For these models some complexity classes indexed by oracles may also be defined. Extension of these classes to probabilistic setting (at least for local decision classes) may be also considered. The goal is here to define and obtain hierarchies similar to classical complexity cases.

In classical computability and complexity, non-determinism plays a fundamental role. The main purpose of Task 5 is to study non-determinism in distributed computing. For example, non deterministic classes of complexity may be defined from the deterministic ones for distributed computing the same way the non-deterministic polynomial time class is defined from the polynomial deterministic time class. Hence, corresponding to the LD class (distributed languages that can be decided in constant time in the $\mathcal{LOCAL}$ model) we can consider the NLD, analogue non-deterministic class, for which there exists a certificate that can be verified in constant number of rounds in the $\mathcal{LOCAL}$ model. Note that as for the complexity classes and "yes-no" problems, several definitions are conceivable. Of course some questions about comparison between deterministic complexity classes and their non-deterministic versions arise naturally.

At last, to complement this classical approach, this project envisions new computational paradigms for theory of distributed computing. We have identified two new paradigms relevant for distributed computing: distributed quantum computing and algorithmic game theory. Quantum-mechanical effects may have an interesting impact on the complexity and especially on its locality aspects for distributed computing [21]. On the other hand, it is interesting to take into consideration game theory aspects in distributed computing.

[4] We prefer the term "yes-no" problem to "decision problem", because "decision problem" has not the same meaning in wait-free or fault-tolerant context.

Some classical problems, specifying the utilities of agents may be reformulated as equilibria in the sense of game theory.

Clearly all these objectives are very challenging. For many of them, the current state of art is still rather primitive. Nevertheless, for most of them a lot of results already exist but generally are not stated in a general framework of complexity theory for distributed computing. Hence we think that most of the objectives are reachable and many new results and new insights into complexity theory for distributed computing are expected from this project.

A crucial aspect of this proposal is that all objectives relate to both communities of Distributed Computing. Most of the time both communities adopt different approaches to tackle those issues. We expect the confrontation of these points of view to be valuable and to enable to converge to an outline of a general complexity and computability theory for distributed computing.

3.2 Description by task

3.2.1 Task 1: Project Management

Coordinator: Carole Delporte-Gallet et Hugues Fauconnier (Paris)
Local coordinators: David Ilcinkas (Bordeaux) et Achour Mostéfaoui (Rennes)

The success of the project is crucial to an intensive collaboration between the three sites, involving frequent meetings and inter-site visits. DISPLEXITY is therefore planing four 2-day meetings every year during which all partners will present their most recent results related to the project. One of these meetings will be dedicated to the compilation of the results achieved during the year, to make sure that all partners synchronize. All dates for deliverables will fit with the dates of these annual special meetings.

In addition, our budget requests a grant for a 1-week visit per year per permanent member (that can be used by the member himself or by his students). The coordinators of the project will pay a specific attention that frequent exchanges between the partners are done.

The management of the web page of the project will play an important role. Indeed, the fundamental nature of the project requires that the results will be made available to the community as quickly as possible. Our web page will be a repository for preliminary versions of papers by the partners, enabling rapid diffusion of our preliminary results inside the project (via a private web site), and a rapid advertising of the finalized results outside the project (via a public web site).

Scientific tasks

Since the project is oriented toward fundamental research, the quality of the project will be evaluated by the ability of the partners to present and advertise their results in the most appropriate and prestigious conferences of the field.

It should be noted that following the philosophy of the project, all the partners are included in all the tasks.

3.2.2 Task 2: "yes-no" and Decision Problems in Distributed Computing

Coordinator: Corentin Travers (Bordeaux)

The overall objective of this task is the study of different settings for tackling decision problems in the various frameworks of distributed computing. As discussed in the general description of the project, we will mostly focus on a notion of decision defined as follows. Given a specification $\mathcal{L}$ (a specification can be formalized as a language-membership), *deciding* $\mathcal{L}$ is defined by the following two requirements:

- If the inputs of the nodes satisfy $\mathcal{L}$ then all nodes must output "yes";
- Otherwise at least one node must output "no".

Hence, the "yes" answer must be unanimous. This definition is perfectly appropriated to, e.g., contexts in which any imperfection should result in a node raising an alarm [24]. On the other hand, one could strengthen the notion of decision by requiring all nodes to output "no" in case of an input not satisfying the given specification $\mathcal{L}$. In other words, one could foresee settings in which both the "yes" and the "no" answers should be unanimous. While such a definition may not be appropriated in frameworks like $\mathcal{LOCAL}$ computation (because a 2-sided unanimous answer may yield communications between far away nodes) or wait-free computation (because consensus is not possible in this model [16]), there are frameworks for which such a definition is well suited. One typical example is mobile agent computing. Indeed, in this setting, it appears that any language that can be decided unanimously 1-sided can also be decided unanimously 2-sided [20].

In fact, one could define decision problems in many different manners. To give just a few examples, consider the following definitions:

- Majority: If the inputs to the nodes satisfy $\mathcal{L}$ then a majority of nodes must output "yes", otherwise a majority of nodes must output "no".
- Weighted: Every node i output a value $x_i \in [0,1]$ so that if the inputs to the nodes satisfy $\mathcal{L}$ then $\sum_i x_i \geq \alpha$, else $\sum_i x_i < \alpha$ where α is some threshold that may depend on the size of the distributed system.
- Functional: Every node i outputs a value x_i belonging to some universe U so that if the inputs to the nodes satisfy $\mathcal{L}$ then $f(x_1,\ldots,x_n) \geq 0$, else $f(x_1,\ldots,x_n) < 0$ where $f : U^* \to \mathbb{R}$.

Objectives of the task:

Sub-Task 2.1: Studying the pertinence of every decision model as a function of the computational model. The main computational models considered in this project are: the wait-free model, the $\mathcal{LOCAL}$ model, the $\mathcal{CONGEST}$ model, the population protocol model and the mobile agent model. One objective on Task 2 is, for each computational model, to determine the most appropriate decision model(s). The central question is: Is there a "natural" universal decision model?

Sub-Task 2.2: Studying the respective powers of the decision models, for a fixed computational model. For instance, for any computational model, a language that can be decided in the 2-sided unanimous model can also be decided in the 1-sided unanimous-"yes" model. More generally, given two functions f and f', and a fixed computational model, how to compare the ability of deciding according to f and the ability of deciding according to f'?

Sub-Task 2.3: Studying the relationship between the decision version of a problem and the computation version of the same problem. For instance, verifying the validity of a coloring can be achieved in one round in the $\mathcal{LOCAL}$ model, whereas computing a valid coloring requires a non-constant number of rounds. Preliminary studies indicate that computing is not always harder than verifying, at least in non-deterministic models. In particular, it is known that verification is sometimes harder than computation in the $\mathcal{CONGEST}$ model and in fault tolerant model. For example the verification of a consensus needs a stronger failure detector than the computation of the consensus. One objective on Task 2 is to tackle the relationship between computing and verifying under various computational models, and for various decision models.

3.2.3 Task 3: Oracles

Coordinator: Michel Raynal (Rennes)

The main objective of this task is to evaluate the possibility and impossibility of what can be computed or more generally achieved in distributed systems. In a rather similar way to classical computation theory, the use of *oracles* is very useful to this end.

Due to the impossibility result of FLP, oracles play a fundamental role concerning fault tolerance: as consensus is not possible in asynchronous models even with only one faulty process, in a very similar way to classical computability, it is natural to add oracles. Among other things, these oracles enable to establish hierarchies between (impossible) problems in distributed fault-tolerant computing. In this way, failure detectors are distributed oracles giving processes information about failures. The formal definition of failure detectors ensures that these oracles depend only on the failures and hence failure detectors encapsulate information about process crashes needed to solve the problem. One important point is that failure detectors may be compared by reduction: failure detector f is stronger than failure detector g if there is a distributed algorithm (a reduction) that enables to output g from f. In this way, failure detectors may be used to compare the difficulty of problems and enable to establish hierarchy between problems that are impossible to solve with process crashes.

Another kind of oracles appears also in wait-free computations. Atomic objects like "test&set", "cmp&swap" are not wait-free implementable in read-write restricted shared memory. In this way, calls to these objects may be considered as calls to oracles. Here, contrary to failure detectors that depend only on failures, the calls of these atomic objects depend also on the specification of the object. One of the main results about wait-free computation is the Herlihy's consensus number [22] that enables to define a hierarchy between computational power of atomic objects.

An important tool in wait-free computation is the BG simulation [4, 5] that enables to simulate systems t-resilient of n processes with systems of wait-free $t + 1$ processes. A very important question is how that kind of simulation may be extended to systems with oracles like failure detectors.

More recently, another approach has been proposed([12]). Instead of having an oracle that gives information about failures, it is possible to consider that an adversary may choose which sets of processes are correct. For example, if we know that only processes a, b, c or a, b or a, c may be the processes that do not crash (the adversary may only choose one of these three sets) it is possible to have an algorithm solving consensus (intuitively a being correct in all the cases, a may realize the consensus). It is then possible to determine for a given problem some structural properties on the set of sets of correct processes that enable to solve this problem. For example, for the k-set agreement task the set of sets of correct processes for which there is a solution roughly corresponds to set of sets of corrects processes having at least $k+1$ correct processes. A very interesting point is to study the relationship between failure detectors and the adversaries. In particular one of the main result concerning failure detectors, define the minimum failure detector to solve the k-set agreement, may be obtained by considering adversaries that enable to solve the k-set agreement.

Concerning the more "network computing" oriented approach, oracles are also interesting. Some very simple oracles may also change the possibility and impossibility results. For example in an anonymous ring of processus, an oracle giving the number of nodes enables to separate rings of different sizes. More generally, the notion of oracles in this context may be an important tool concerning complexity issues.

From a more practical point of view, many problems in distributed computing lead to graph representations where oracles may have an important role. For example, a peer to peer system is a graph, representing the topology of the logical network. Many relationships between entities can be represented by an edge in a graph: for instance a social network can be represented by a graph where there is an edge between 2 users if they are friends on Facebook, or an edge could be present between an item and a node if the node has used or tagged or bought an item (for example recommenders systems rely on such graphs). There are two main classes of problems that emerge: *(i)* either a graph exists and the goal is to extract some information from the graph that can then be used to provide a given functionnality in a distributed system (this includes studying the graph properteis of a social graph for example), or *(ii)* a graph is created to achieve some functionality of a system and its properties may help to tune the system (for example creation of a given overlay and study the properties of the resulting graph). This is typically some problems that should concern both communities.

Objectives of the task:

Sub-Task 3.1: Define a general theory of oracles in distributed computation is a preliminary goal of this task. For this the notion of reduction between oracles is central. The extension of usual reductions and simulations between systems to reductions and simulations of systems with oracles is another important goal of this task.

Sub-Task 3.2: Concerning wait-free or fault-tolerant computation, the relationship between adversaries and oracles would enable to obtain new results. Some of these results may have practical impact and would help to refine the hierarchy between impossible problems.

Sub-Task 3.3: An objective is to clearly define what are the graph properties expressed as oracles that are important in the context of distributed computing and to what extent they can help to improve the system. An another objective is to be able to capture such properties such the graph size, the community structure, the bottlenecks (for instance provided by centrality measure), etc.

Sub-Task 3.4: Finally the previous objectives mentioned should be considered in the context of a dynamic and large-scale system, clearly leading to approximation. Calibrating such approximations and studying the impact on their utility is of crucial importance.

3.2.4 Task 4: Complexity classes

Coordinator: Pierre Fraigniaud (Paris)

The overall objective of this task is to design significant complexity classes in the different frameworks of distributed computing. This design has to be made having in mind the general "wish" of eventually being able to compare classes defined for *different* frameworks, e.g., to be able to compare a class defined for the wait-free model with a class defined for the $\mathcal{LOCAL}$ model. The definition of the complexity classes is subject to the notion of decidability one is considering. We foresee that the 1-sided unanimous "yes" notion is the most promising (cf. Task 2), and thus Task 4 is described according to this assumption. Nevertheless, it may be the case that Task 4 should be revisited if it occurs that our investigations performed in Task 2 demonstrate that other notions of decidability should also be considered.

The main basic complexity classes that will be investigated in the project are the following:

- LD (for Local Decision) is the class of distributed languages that can be decided in constant time in the $\mathcal{LOCAL}$ model;
- LDC is the class of distributed languages that can be decided in constant time in the $\mathcal{CONGEST}$ model;
- WFD (for Wait-Free Decision) is the class of distributed languages that can be decided in the wait-free model;
- MAD (for Mobile Agent Decision) is the class of distributed languages that can be decided in the mobile agent model;
- POPD (for Population Protocol Decision) is the class of distributed languages that can be decided in the population protocol model.

One note a significant difference between, on the one hand, the classes LD and LDC, and, on the other hand, the classes WFD, MAD and POPD. The former refers to a notion of time (number of rounds in the $\mathcal{LOCAL}$ or in $\mathcal{CONGEST}$ model). They have thus the flavor of computational complexity. The latter instead does not refer to any efficiency measure, and focuses solely on the ability of achieving decision. They have thus a flavor of decidability classes. Nevertheless, as it will be described in Task 5, there seems to be connections between these two very different types of classes, at least in the case of non-deterministic computation.

It is also worth noticing that these complexity classes can be generalized to a probabilistic setting. For instance, one can define the class BPLD (for Bounded-error Probability Local Decision) as the set $\cup_{0\leq p,q\leq 1}\text{BPLD}(p,q)$ where $\text{BPLD}(p,q)$ is the class of all distributed languages that can be decided by a randomized distributed algorithm that runs in a constant number of communication rounds (in the $\mathcal{LOCAL}$ model) and produces correct answers on legal (respectively, illegal) instances with probability at least p (resp., q).

Objectives of the task:

Sub-Task 4.1: Characterizing the complexity classes listed above, or at least some fragments of them. For instance, it can be shown that by slightly reducing the definition, the languages in the class WFD are in one-to-one correspondence with covering spaces, in the sense of algebraic topology. In general, one objective of this sub-task is to determine whether these classes are decidable. Defining pertinent complexity measures for the wait-free model as well as for the mobile agent model is another objective. Notions like number of memory accesses, number of interactions or number of moves appear to be natural for the wait-free, population protocol and mobile agent model, respectively.

Sub-Task 4.2: Studying the impact of randomization on the deterministic complexity classes LD, LDC, WFD, and MAD. For instance, it has been recently shown [19] that LD and BPLD coincide for a large range of values p and q, if one restricts the languages to the hereditary ones. The question "does randomization help?" is at the core of computer science. Proving or disproving LD = BPLD is probably beyond the objectives of this project. Nevertheless, we will work having this type of question in mind.

Sub-Task 4.3: Defining appropriate notion of reductions that enable to compare the "difficulty" of problems belonging to the same class. In the $\mathcal{LOCAL}$ model, a straightforward notion of reduction is to consider local computation mapping one instance to another while preserving the respective language membership. Defining the analog in, e.g., the framework of wait-free computation appears to be more challenging, because the notion of termination is more difficult to handle, and because faults can occur during the reduction.

3.2.5 Task 5: Non-determinism in distributed computing

Coordinator: Amos Korman (Paris)

One can consider LD as the analogue of the class P, and then define the class NLD as the analogue of the class NP. Specifically, NLD (for non-deterministic local decision) is the class of distributed languages for which there exists a *certificate* (or *proof*) that can be verified in a constant number of communication rounds in the $\mathcal{LOCAL}$ model. More precisely, NLD is the class of distributed languages $\mathcal{L}$ for which there exists a local algorithm such that:

- if $L \in \mathcal{L}$, then there exists a certificate c such that algorithm A applied on L with certificate c outputs "yes" for all nodes,
- otherwise, for every certificate c, algorithm A applied on L with certificate c outputs "no" for at least one node.

It is worth noticing that non-deterministic distributed computing does not find its interests in complexity theory only. In particular, the above notion of non-determinism is very much related to the area of computation with advice (e.g., [17, 18]), which aims at investigating the amount of information known to each node and its impact on the performances of the algorithm. In fact, the definition of NLD finds most of its similarities with the notion of proof labeling schemes (cf., e.g.,[23, 24]) which finds applications in the context of fault-tolerance and self-stabilization. Roughly, in certain contexts, one can afford spending lot of time computing the certificate, but once it is given to the nodes, one wants the verification to be fast. (One example is fault-detection in environments such as a plane or a nuclear plant). Clearly, LD $\subseteq$ NLD, and it can be shown that LD $\neq$ NLD. Interestingly enough, it can be also shown that there exists a problem which is NLD-complete for the local reduction mentioned in Task 4. Still, the nature of NLD is far from being well-understood.

Similarly, one can define non-deterministic versions of the classes LDC, WFD, and MAD. In the case of the mobile agent model, restricted to a single agent, it is form instance known that NMAD $\cup$ co-NMAD = MAD. As we will see, non-determinism finds all its interests when combined to the notion of oracles, which will be the objective of Task 3. Task 4 focuses on non-determinism per se.

Objectives of the task:

Sub-Task 5.1: One of our main concerns will be to identify the languages in NLD and NMAD. In particular, from a purely graph theoretic point of view, it is an intriguing question to identify which graph families belong to NLD. We are also aiming at identify subclasses of NLD defined according to the certificate size, finding connections between these subclasses, and finding complete problems for each of the subclasses.

Sub-Task 5.2: Finding trade-offs between the certificate size and the locality of the verification. Indeed, the amount and type of information provided to each node is a central issue in the design of distributed algorithms and local algorithms in particular. It is however interesting to point out that most applications for proof labeling schemes do not perform in constant time. Thus, from the point of view of these applications, there is no real need to restrict the running time of the proof labeling schemes to constant. Reducing the certificate sizes of known proof labeling schemes at the price of somewhat increasing the locality of the verification is a challenging task that may thus contribute to the design of more compact applications. (For instance, doing this in the context of MST verification may yield breakthroughs in the strives for finding an efficient self-stabilizing MST construction protocol that is optimal in terms of memory).

Sub-Task 5.3: Understanding the apparent lack of impact of non-determinism in the framework of wait-free computing. In other words, understanding why non-determinism does not seem to help in the context of wait-free computations. On the other hand, we have pointed out above that WFD can be essentially described in term of universal covers. It has been recently observed that output-free languages in NLD have strong connections with graph coverings. One objective of Task 5 will be to investigate the potential intriguing connections between WFD and NLD, which may not be a coincidence. Needless to say that establishing connections between classes defined for different distributed computational model would be a significant achievement.

3.2.6 Task 6: New computational paradigms/frameworks

Coordinator: Cyril Gavoille (Bordeaux)

The principal objective of this task is to discover to what extent it is possible to improve the distributed complexity of combinatorial optimization problems by taking advantage of aspects of "new" computational paradigms. Inspired by recent developments of computational complexity, we propose to focus on two aspects: quantum-mechanical effects, and algorithmic game theory aspects.

Our studies will involve different models of distributed computing, including processor network graphs, mobile-agent-based computing, and tasks involving the distribution of topological information (computing distance labelings, compact routing tables).

The introduction of computational models based on quantum computing, starting from the works of Deutsch in the 1980's [14], has led to the advent of a new branch of complexity theory. Independently of this, properties of quantum-mechanical systems have proven to be of interest from the perspective of game theory, information theory, and distributed systems [6, 13]. A major line of study concerns the application of quantum entanglement to reduce communication complexity, i.e., to decrease the number of communication bits required to solve a specific task performed within a system graph with several distributed agents [9]. The influence of quantum information on the computing power of distributed systems with node anonymity and distributed systems in the presence of faults has also been studied for problems such as leader election or distributed consensus [28]. The goal of our project is to extend these results and obtain new insight which would be applicable in the context of distributed computing, especially in a combinatorial setting.

In an orthogonal way, recent computational complexity have shown the necessity to revisit classical distributed algorithmic with the eyes of game theory. When conceiving a distributed system, it not sufficient to ask to all involved agents to act according to some algorithm, but also, as agents often correspond to various agents with their own economical aspects, to guarantee that it is indeed in their economical interest to do so. This often leads to add some game theoretic model to the system, specifying the utilities of involved agents, and then revisiting whether situations computed globally correspond to equilibria in the sense of game theory.

Objectives of the task:

Sub-Task 6.1: *Reducing the time complexity of distributed algorithms using quantum effects.* The starting point for our considerations is the well-established $\mathcal{LOCAL}$ model of distributed computing with a network graph of processors exchanging messages of unbounded size. In this context, we wish to discover how the introduction of a globally pre-entangled state of the system or the application of quantum communication channels can decrease the number of rounds required to solve classical combinatorial optimization problems, such as graph $(\Delta + 1)$-coloring. Currently, only some very simple proof-of-concept examples of such problems are known [21]. We are especially interested in effects which require quantum entanglement and cannot be obtained using purely classical correlations, such as shared random bits.

Sub-Task 6.2: *Providing a comparison of the "computational power" of the quantum and non-quantum models*, formalising the notion of locality in quantum distributed computing, and showing how it essentially differs from the understanding of locality in the $\mathcal{LOCAL}$ model. In particular, we would like to obtain more precise algorithmic characterizations of probability distributions of feasible outputs which can be obtained by a quantum system in a given number of rounds. We also intend to elaborate lower bounds on the complexity of problems in quantum models by making use of more powerful models of computations based on stronger-than-quantum non-local boxes (Popescu-Rohrlich boxes and their variants), which we intend to define for the purpose.

Sub-Task 6.3: *Identifying the potential of quantum information in topological queries.* We would like to study the potential of quantum information in reducing the size of distributed data structures required for tasks such as adjacency labeling or distance labeling of a graph, allowing the recovery of information about the location of a pair (or more generally subset) of nodes of a graph based only on the information stored within these nodes. In this sense, we would like to establish if information "spread out" over the whole graph using entangled states can prove more efficient in local queries than classical information (without globally violating the information bound implied by Holevo's theorem).

Sub-Task 6.3: *Providing a comparison of the "computational power" of models when rationality is imposed and when it is not*, formalising when local rules in distributed computing correspond to rational rules in the sense of game theory, and showing how it may or may not influence the computational power of the involved models. In particular, we would like to understand how complexity classes proposed in the other tasks differ or not from their analog when the hypothesis of rationaly is imposed to agents. This involves to consider the various possibilities to go from a game to a distributed system (ex. by repeating local games, by evolutionay game theory models) and then characterize what can be computed by which model using the corresponding notion of rationality.

3.3 Tasks schedule, deliverables and milestones

All the tasks listed in the previous sections are in symbiosis, and there is no precedence constraints among them. Each of the tasks is related to a well identified problem, and the resolution of any problem will benefit to many others. In fact, the tasks are as many different ways of tackling key fundamental problems in the framework of distributed computing. Moreover, because of the theoretical research flavor of the tasks, it is quite hard (and probably inappropriate) to decide when a task should start and end.

On the other hand, DISPLEXITY will pay attention to allow the participants to the project to be cross fertilized by the results of all its members. This will be implemented via frequent meetings and visits between the three sites. By doing so, the participants to DISPLEXITY will share a common expertise, and be able to solve more and more complex problems related to the tasks of the project. Within the 4-years

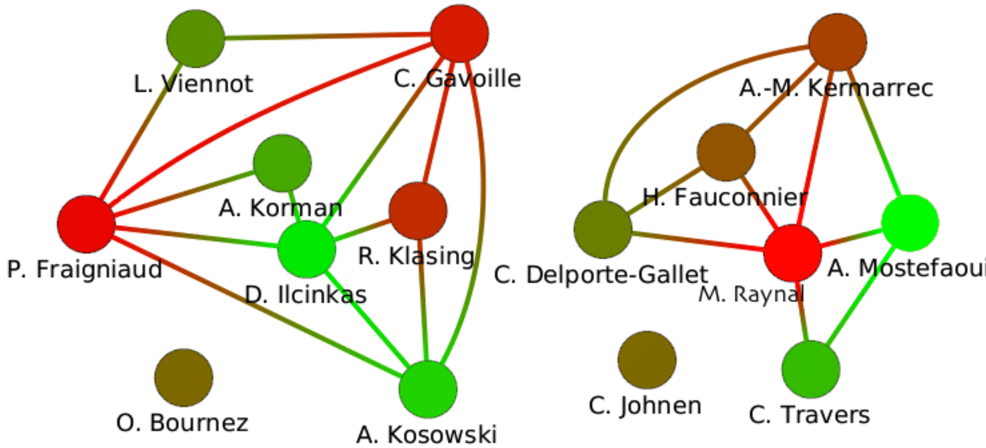

Figure 1: Interaction graph beween partners (DBLP).

duration of the project, it is therefore expected that many key problems will have been solved, providing as many breakthroughs in the field of distributed computing.

One of these meetings will be dedicated to the compilation of the results achieved during the year. The dates for deliverables (t+12, t+24, t+36, t+48) containing the progress of each task will fit with the dates of these annual special meetings.

4 Dissemination and exploitation of results, intellectuel property

The main expected results of this project are scientific advances in the investigated research fields. Research results will be presented at the leading international conferences and in the leading scientific journals.

Each year, PODC, the top level ACM conference on distributed computing invites proposals for workshops that are closely related to the scope of the conferences, i.e. on topics related to the theory, design, specification or implementation of distributed systems. Workshops are in conjunction with the conference on the day immediately preceding or following the conference program. The organization of such a workshop will be an opportunity to promote our scientific results and to draw the attention of the international community. And we plan to propose such a workshop the third year of his project.

DISPLEXITY will maintain a project Web Site that will make available the project results; it will be a useful portal open to all researchers for accessing papers and deliverables.

5 Consortium description

5.1 Partners description and relevance, complementarity

The consortium gathers members from three French leading labs in Distributed Computing (LIAFA, LaBRI, and IRISA). Each of these labs has a worldwide reputation of excellence in the design and analysis of distributed algorithms. DISPLEXITY has paid attention to include experts of most sub-topics in distributed computing, including network computing, mobile computing, fault-tolerant computing, population protocols, failure detector theory, etc. In order to illustrate the high level of expertise of the DISPLEXITY partners, let us just mention that the consortium includes several ACM PODC and DISC Program Committee Chairs (which are the top Int'l conferences in the domain), several DISC and SIROCCO Steering Committee Chairs, and regular Program Committee members of most conferences in distributed computing (ACM PODC, DISC, SIROCCO, OPODIS, SSS, IEEE ICDCS, IPDPS, etc.) and Editorial Board members of top Int'l journals (e.g., IEEE Transactions on Computers, IEEE Transactions on Parallel and Distributed Systems, Journal of Parallel and Distributed Computing).

The quality of the consortium is more than the sum of the quality of its members. Indeed, the main interest of DISPLEXITY will be its ability to tackle the main research models in distributed computing in parallel. To enable DISPLEXITY to take up the challenges of dealing with apparently so many different computational models, we have focused our attention on yes/no problems, which will ensure a global coherency of the project. This apparent narrowing of the problem field is actually not restrictive as it is sufficient for the development of a complexity theory (this is because such theory is based on decision problems).

As illustrated by the collaborative diagram displayed on Figure 1, the consortium includes two well identified sub-groups, with each of the sub-groups scattered over the three collaborative institutions LIAFA-LaBRI and IRISA. Roughly, one sub-group includes individuals dealing with problems related to asynchrony and faults, while the other one includes partners dealing with network computing.

It should be mentioned that many participants regularly attend to the same conferences gathering both communities. The high added value brought by DISPLEXITY is to provide a general framework (Computational Complexity) in which these people will exchange and collaborate. Conversely, the richness of the consortium, beyond the individual quality of its members, is to gather these two groups together: each one separately will not be able to cover more than a fragment of the theory. We do hope, and actually strongly believe, that the same drawing at the end of DISPLEXITY will display a single component, with connecting edges corresponding to many significative publications in Distributed Computational Complexity.

Bordeaux (LABRI): The LaBRI (Laboratoire Bordelais de Recherche en Informatique) is a research unit associated with the CNRS (UMR 5800), the University of Bordeaux 1, the IPB and the University of Bordeaux 2. It has significantly increased in staff numbers over recent years and now includes around 350 members (academics, researchers, PhDs, etc.). Today the members of the laboratory are grouped in six teams, each one combining basic research, applied research and technology transfer : * Combinatorics and Algorithmics * Image and Sound * Languages, Systems and Networks * Formal Methods * Models and Algorithms for Bio-informatics and Data Visualisation * Supports and Algorithms for High Performance Numerical Applications

The LaBRI participants to DISPLEXITY are all members of the Combinatorics and Algorithmics team and most of them are in the INRIA team CEPAGE. They are all recognized worldwide experts in distributed computing. The LaBRI team includes specialists on various models and topics in the scope of the project, in particular fault tolerance issues, network and communication algorithmics, and (possibly quantum) distributed information.

Paris (LIAFA): The LIAFA (Laboratoire d'Informatique Algorithmique: Fondements et Applications) is supported jointly by the French National Center for Scientific Research (CNRS) and by the University Paris Diderot - Paris 7. It is member of the Fondation Sciences mathématiques de Paris. The main research topics addressed by LIAFA are related to theoretical computer science. The one hundred members of LIAFA (academics, researchers, and PhDs) are divided in five research teams: Algorithms and Complexity, Combinatorics, Distributed Algorithms and Graphs, Automata and applications, and Modeling and verification.

DISPLEXITY involves all members of the team Distributed Algorithms and Graphs whose research interests are dealing with distributed computing. All these people are worldwide experts in distributed computing, covering most of the topics in the field, both theoretically-oriented and practically-oriented (they are all members of the INRIA team-projet ”GANG”).

Rennes (IRISA-ASAP): The ASAP (As Scalable As Possible) team, lead by Anne-Marie Kermarrec, is part of IRISA/INRIA Rennes a research unit associated mainly with University of Rennes, CNRS and INRIA. The research activities of ASAP range from theoretical foundations to practical protocols and implementations for (mainly large-scale and dynamic) distributed systems in order to cope with the recent and tremendous evolution of distributed systems. Effectively we observed huge evolutions: (i) Scale shift in terms of system size, geographical spread, volume of data, and (ii) Dynamic behaviour due to versatility, mobility, connectivity patterns.

The research of the ASAP Project-Team is along two main themes: *Distributed computing models and abstractions* and *Peer-to-peer distributed systems and applications*. These research activities encompass both long term fundamental research seeking significant conceptual advances, and more applied research to validate the proposed concepts against real applications. ASAP regroups 14 researchers (including one professor, one senior researcher, two associate professors and one researcher).

The proposal also involves one researcher from team Algorithmic and Complexity from Computer Science Laboratory (LIX) from Ecole Polytechnique. The research activity of this member is engaged in research along two themes related to this proposal: models of computation, and in particular distributed models of computation and algorithmic game theory.

5.2 Qualification of the proposal coordinator

Carole Delporte-Gallet and Hugues Fauconnier will share the coordination task. They have a long-standing collaboration in research projects and have a strong experience as investigators of projects and organizators of workshops on Distributed Computing. They are top level experts in Fault Tolerance Distributed Computing and they published papers in many journals like JACM, Distributed Computing, TOPLAS and in top level conferences of this aera (PODC, DISC, DSN, ICDCS, ...). They have been program committee members of established conferences in Distributed Computing such as PODC, DISC, IEEEE ICDCS, OPODIS, SIROCCO...

As investigators of projects, they have been investigators in the last past years of the Action Spécifique: Algorithmes Distribués (03-04) and of the BQR-action “Distributed Quantum computing”of the University Paris-Diderot University’ (2010). They organized a school on Distributed Computing in 2003 and two workshops: “Algorithmique distribuée et applications” in 2004 and ”Distributed Quantum computing” in 2010.

C. Delporte-Gallet is Professor in computer science at the University Paris-Diderot and leader of the team “Distributed Algorithms andGraphs” of the LIAFA . She was alumni of the ENS Sèvres (79-83). She holds her Ph.D. in Computer Science from the University Paris Diderot in 1983 and is HDR in 2001.

H. Fauconnier received his Ph.D. in 1982 and HDR degree in 2001 in Computer Science from the University Paris-Diderot, after Master degrees in Mathematics and Computer Science. He is maitre de conferences in Computer Science at the University Paris-Diderot.

Hugues Fauconnier was local investigator for the ACI project FRAGILE (2006-2008) and for the ANR-VERSO SHAMAN (2008-2012). It is also local investigator of an Ile de France Post-Doc program PEFI-CAMO (2009-2011).

5.3 Qualification and contribution of each partner

Partner 1: Paris(Coordinator)

Name	First name	Position	Field of research	PM	Contribution to the proposal (4 lignes max)
Delporte	Carole	PR	Distributed computing, fault tolerance (t-resilent, wait free, failure detector, adversaries), population protocol	36	Coordinator of the project(Task1)
Fauconnier	Hugues	MCF	Distributed computing, fault tolerance (t-resilent, wait free, failure detector, adversaries), population protocol	36	Coordinator of the project (Task1)
Fraigniaud	Pierre	DR	Distributed computing, Network computing, Mobile computing	36	Coordinator Task 4
Korman	Amos	CR	Distributed computing, labeling schemes, sparse spannners	36	Coordinator Task 5
Viennot	Laurent	DR	Networks	12	
Bauman	Hervé	PhD	Rumor spreading	12	
Koegler	Xavier	PhD	Population protocol	24	
Arfaoui	Heger	PhD	Distributed computing complexity in the $\mathcal{LOCAL}$ model	36	

Partner 2: Rennes

Name	First name	Position	Field of research	PM	Contribution to the proposal (4 lignes max)
Mostefaoui	Achour	MCF	Distributed computing. Implementation of Oracles and design of data structures well suited to a given oracle	36	local coordinator (Task 1)
Raynal	Michel	PR	Distributed computing. Decidabiliy of decision tasks and lower bounds for implementation in a given system	36	coordinator Task 3
Kermarrec	Anne-Marie	DR	Peer-to-peer Systems. Definition and implementation of oracles well-suited to peer-to-peer systems	12	
Bournez	Olivier	PR	Computability and complexity with distributed models, anonymous models, analog models. Algorithmic game theory.	24	
Imbs	Damien	Phd Student	Distributed computing. Decidabiliy of decision tasks and lower bounds for implementation in a given system	24	

Partner 3: LABRI

Name	First name	Position	Field of research	PM	Contribution to the proposal (4 lignes max)
Gavoille	Cyril	PR	Distributed Graph Algorithms, Quantum Information	36	Task 6 Coordinator
Ilcinkas	David	CR	Distributed Graph Algorithms, Mobile Agent Computing	36	Local coordinator (Task 1)
Johnen	Colette	PR	Fault Tolerance	28	
Klasing	Ralf	DR	Network algorithms, Mobile Agent Computing	24	
Kosowski	Adrian	CR	Distributed Comptuting, Quantum Information	20	
Travers	Corentin	MCF	Fault Tolerance	36	Task 2 Coordinator
Halftermeyer	Pierre	PhD	Disributed Distance Oracles	12	
Wade	Ahmed	PhD	Mobile Agent Computing, Dynamic Graphs	24	

6 Scientific justification of requested ressources

The budget concerning small material will mostly be used for providing participants with appropriate computers, and to give adequate resources to the PhDs and Post-doc involved in the project.

A large part of the requested grant concerns missions. For the success of the project, it is indeed crucial that the participants get enough supports for frequent exchanges and visits between partners. In addition, the topics addressed by the project are quite hot, and evolve rapidly. It is thus very important that the participants of DISPLEXITY can participate to the most relevant conferences of the field, for presenting their results, and to carry on collaborations with colleagues. The yearly budget for missions is actually increasing with time, it is planed that the need for traveling will increase with time, for participants are expected to present their new results in various places, including the most visible scientific events.

Most of the partners of this project plan to concentrate their research on DISPLEXITY during the four next years.

6.1 Partner 1: Paris

	Amount
Personnel	
PhD	3 years
Master Internship	4* 6 months
Mission	178 000
Small equipment	20 000

- Equipment

  None.

- Staff

  PhD LIAFA: "yes-no decision problems"

  This PhD position is related to Task 2 ("yes-no" and Decision Problems in Distributed Computing). Its first objective is, for the wait-free model and the $\mathcal{LOCAL}$ computational model, to determine the most appropriate decision model. The second objective will be to study the respective powers of the decision model(s), for a fixed computational model and the relationship between the decision version of a problem and the computation version of the same problem. For instance, verifying the validity of a coloring can be achieved in one round in the $\mathcal{LOCAL}$ model, whereas computing a valid coloring requires a non-constant number of rounds. For example the verification of a consensus needs a stronger failure detector than the computation of the consensus. One objective of Task 2 is to tackle the relationship between computing and verifying under various computational models, and for various decision models.

  This PhD position will be open to the best Master students in France and the aforementioned foreign universities.

  Master student internship Master student internship will be proposed in relation with the different aspects of the project and more precis task 3, 4, 5 and 6. It is expected to supervise one master student per year for a six months period each (the typical internship length in our university).

- Subcontracting None

- Mission

  The total amount for travel is 178 000 and is justified as follows:

  - As pointed out in the Task 1, four 2-days workshops will be organized each year which induce a cost of 4 500 Euros per such meeting for all members. The resulting cost is 76 000.

- As pointed out in the Task 1, our budget requests a grant for a 1-week visit per year per permanent member for a total cost of 20 000 Euros.
- It is expected that most of the results will be communicated to top level international conferences in distributed computing. Moreover, as explained at the beginning, it is also vitally important that the members collaborate on the project topic with the other (international) experts on the domain. The resulting cost is roughly estimated to 82000 Euros. It could be considered as rather high, but it should be taken into account that the permanent members of the Paris partner published 82 papers in the major conferences of the domain in the last four years (the same duration as the project). Even without taking into consideration the collaborations with other colleagues, the participation to major workshops without refereed publications, and the missions of the non permanent staff, this could already almost justify the value of 82 000 Euros for missions.

• Costs justified by internal invoicies

• Other expenses

The budget concerning small material will mostly be used for providing participants with appropriate computers, and to give adequate resources to the PhDs and Post-doc involved in the project. Various expenditures will be made throughout the project duration for a total cost of 20 000 Euros.

6.2 Partner 2: Rennes

	Amount
Personnel	
Post-doc	18 months
Master Internship	4*6 months
Mission	112 000 E
Small equipment	10 000 E

• Equipment

None.

• Staff

The post-doc student will be hired for working mainly on Task 3 devoted to Oracles. Of course, she will also participate to the other tasks. The notion of oracle has been introduced many years ago mainly for solving the consensus problem is asynchronous systems but has been extended to tackle many other decision problems and among others we propose to extend this notion to the graph topology that underlies any computing network.

The master student internships will be proposed to students in order to work on specific aspects on the design of distributed data structures that rely on some oracle. These aspects will be connected to tasks 5, 6 and 3. The internships will be proposed on a one or two per year basis probably starting from the second year of the project.

• Subcontracting

None

• Mission

The total amount for travel is 112000 euros and is justified as follows:

- As pointed out in the Task 1, our budget requests a grant for a 1-week visit per year per member for a total cost of 20000 Euros.

- As pointed out in the Task 1, four 2-days workshops will be organized each year which will induce a cost of 1 250 Euros per such meeting for all members. The resulting cost is 20 000.
- It is expected that most of the results will be communicated to top level international conferences in distributed computing. Moreover, as explained at the beginning, it is also vitally important that the members collaborate on the project topic with the other (international) experts on the domain. The resulting cost is roughly estimated to 30 000 Euros. It could be considered as rather high, but it should be taken into account that the permanent members of the Rennes partner published more than a hundred papers in the major conferences of the domain in the last four years (the same duration as the project). Even without taking into consideration the collaborations with other colleagues, the participation to major workshops without refereed publications, and the missions of the non permanent staff, this could already almost justify the value of 72 000 Euros for missions (on a basis of two int. confrences per permanent member per year).

- Costs justified by internal invoicies
- Other expenses

  The amount of 10000 Euros for small equipment is for purchasing four laptops. One for the post-doc student, one for the different master students over four years and the two others are for a renewal of machines for two of the members of the team involved in this ANR project.

6.3 Partner 3: Bordeaux

	Amount
Personnel	
Post-doc	1.5 years
Master Internships	20 months
Mission	164000 Euros
Small equipment	15000 Euros

- Equipment

  None.

- Staff

  Post-doc position. The 18-month post doc will support us in Task 4 dedicated to the complexity classes and will particularly focus on studying the impact of randomization on the deterministic complexity classes (Task 4.2). A strong expertise in complexity theory and randomized algorithms will be requested in the position description. The latter will be disseminated worldwide through the main e-mailing list of the theoretical computer science community. This post-doc induces a cost of 1.5*42000 = 63000 Euros.

  Master student internships. Master student internships will be proposed in relation with the different aspects of the project. It is expected to supervise one master student per year for a five-month period each (the typical internship length in our university). This induces a cost of 4*5*417,09 = 8341.80 Euros.

- Subcontracting None.
- Mission

  The total amount for travel is 164000 Euros and is justified as follows:

  - As pointed out in Task 1, four 2-days workshops will be organized each year which induce a cost of 3500 Euros per such meeting for all members. The resulting cost is 56000 Euros.

- As pointed out in Task 1, our budget requests a grant for a 1-week visit per year per permanent member for a total cost of 24000 Euros.
- It is expected that most of the results will be communicated to top level international conferences in distributed computing. Moreover, as explained at the beginning, it is also vitally important that the members collaborate on the project topic with the other (international) experts on the domain. The resulting cost is roughly estimated to 84000 Euros. It could be considered as rather high, but it should be taken into account that the permanent members of the Bordeaux partner published 84 papers in the major conferences of the domain in the last four years (the same duration as the project). Even without taking into consideration the collaborations with other colleagues, the participation to major workshops without refereed publications, and the missions of the non permanent staff, this could already almost justify the value of 84000 Euros for missions.

- Costs justified by internal invoicies

  None.

- Other expenses

  The budget concerning small material will mostly be used for providing participants with appropriate computers, and to give adequate resources to the PhDs and Post-doc involved in the project. Various expenditures will be made throughout the project duration for a total cost of 15000 Euros.

7 Annexes

7.1 References

References

[1] Dana Angluin, James Aspnes, Zoë Diamadi, Michael J. Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, 18(4):235–253, 2006.

[2] Franois Bonnet and Michel Raynal. Anonymous asynchronous systems: The case of failure detectors. In *DISC 2010*, pages 206–220, 2010.

[3] Franois Bonnet and Michel Raynal. Consensus in anonymous distributed systems: Is there a weakest failure detector? In *AINA 2010*, pages 206–213, 2010.

[4] Elizabeth Borowsky and Eli Gafni. A simple algorithmically reasoned characterization of wait-free computations (extended abstract). In *PODC*, pages 189–198, 1997.

[5] Elizabeth Borowsky, Eli Gafni, Nancy A. Lynch, and Sergio Rajsbaum. The bg distributed simulation algorithm. *Distributed Computing*, 14(3):127–146, 2001.

[6] Anne Broadbent and Alain Tapp. Can quantum mechanics help distributed computing? *ACM SIGACT News - Distributed Computing Column*, 39(3):67–76, September 2008.

[7] Tushar Deepak Chandra, Vassos Hadzilacos, and Sam Toueg. The weakest failure detector for solving consensus. *J. ACM*, 43(4):685–722, July 1996.

[8] Tushar Deepak Chandra and Sam Toueg. Unreliable failure detectors for reliable distributed systems. *J. ACM*, 43(2):225–267, March 1996.

[9] Richard Cleve and Harry Buhrman. Substituting quantum entanglement for communication. *Physical Review A*, 56(2):1201–1204, Aug 1997.

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

[11] Carole Delporte-Gallet, Hugues Fauconnier, Rachid Guerraoui, and Anne-Marie Kermarrec. Brief announcement: byzantine agreement with homonyms. In Friedhelm Meyer auf der Heide and Cynthia A. Phillips, editors, *SPAA*, pages 74–75. ACM, 2010.

[12] Carole Delporte-Gallet, Hugues Fauconnier, Rachid Guerraoui, and Andreas Tielmann. The disagreement power of an adversary. *Distributed Computing*, 2010. Accepted in DISC: Special issue.

[13] Vasil S. Denchev and Gopal Pandurangan. Distributed quantum computing: A new frontier in distributed systems or science fiction? *ACM SIGACT News - Distributed Computing Column*, 39(3):77–95, September 2008.

[14] David Deutsch. Quantum theory, the Church-Turing principle and the universal quantum computer. *Proceedings of the Royal Society of London*, A400:97–117, 1985.

[15] Shlomi Dolev. *Self-Stabilization*. MIT Press, 2000.

[16] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, April 1985.

[17] Pierre Fraigniaud, Cyril Gavoille, David Ilcinkas, and Andrzej Pelc. Distributed computing with advice: Information sensitivity of graph coloring. In 34^{th} *International Colloquium on Automata, Languages and Programming (ICALP)*, volume 4596 of Lecture Notes in Computer Science, pages 231–242. Springer, July 2007.

[18] Pierre Fraigniaud, Amos Korman, and Emmanuelle Lebhar. Local mst computation with short advice. In Phillip B. Gibbons and Christian Scheideler, editors, *SPAA*, pages 154–160. ACM, 2007.

[19] Pierre Fraigniaud, Amos Korman, and David Peleg. Local distributed decision. *CoRR*, abs/1011.2152, 2010.

[20] Pierre Fraigniaud and Andrzej Pelc. Decidability classes for mobile agents computing. *CoRR*, abs/1011.2719, 2010.

[21] Cyril Gavoille, Adrian Kosowski, and Marcin Markiewicz. What can be observed locally? round-based models for quantum distributed computing. Technical report, arXiv: quant-ph/0903.1133, 2009.

[22] Maurice P. Herlihy. Wait-free synchronization. *ACM Transactions on Programming Languages and Systems*, 13(1):123–149, January 1991.

[23] Amos Korman and Shay Kutten. Distributed verification of minimum spanning trees. *Distributed Computing*, 20(4):253–266, 2007.

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

[25] Nathan Linial. Locality in distributed graphs algorithms. *SIAM Journal on Computing*, 21(1):193–201, 1992.

[26] Moni Naor and Larry J. Stockmeyer. What can be computed locally? *SIAM J. Comput.*, 24(6):1259–1277, 1995.

[27] David Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM Monographs on Discrete Mathematics and Applications, 2000.

[28] Seiichiro Tani, Hirotada Kobayashi, and Keiji Matsumoto. Exact quantum algorithms for the leader election problem. In 22^{nd} *Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 3404 of Lecture Notes in Computer Science, pages 581–592. Springer, February 2005.

7.2 CV, resume

Partner 1: Paris

DELPORTE-GALLET Carole

Coordinator

Age: 51
Position: PR
Email: delporte@liafa.jussieu.fr
URL: www.lliafa.jussieu.fr/~cd

LIAFA, Université Paris Diderot
Case 7014
F-75205 Paris cedex 13
Tel: (+33) 1 57 27 92 25

Cursus

PhD (1983) and HDR (2001) from University Denis Diderot

Publications

Number of publications in refereed international journals: 10
(*JACM, TOPLAS, Distributed computing, Information and Computation, Information and Control, JPDC...*

Number of publications in refereed international conferences with proceedings: 39
(*PODC, DISC, DSN, ICDCS, ICALP, SPAA, SRDS, OPODIS, SIROCCO, IPDPS, ICDCN, DCOSS....)*

Selected publications from the past five years
(Authors are given by the alphabetical order of their name)

- C.Delporte-Gallet, H.Fauconnier, R.Guerraoui and A. Tielmann. **The Weakest Failure Detector for Message Passing Set-Agreement**. *Distributed computing*, online . (2010)
- M.K. Aguilera,C.Delporte-Gallet, H.Fauconnier and S. Toueg. **Partial synchrony based on set timeliness**. *PODC*, pages 102-110. (2009)
- C.Delporte-Gallet, S. Devismes and H.Fauconnier **Stabilizing leader election in partial synchronous systems with crash failures**. *J. Parallel Distrib. Comput.*, volume 70 (1), pages 45-58,. (2010)
- C.Delporte-Gallet, H.Fauconnier, R.Guerraoui and A. Tielmann. **The Weakest Failure Detector for Message Passing Set-Agreement**. *DISC*, pages 109-120). (2009)
- C.Delporte-Gallet, H.Fauconnier, F. Freiling, L. Penso and A. Tielmann. **From Crash-Stop to Permanent Omission: Automatic Transformation and Weakest Failure Detectors**. *DISC*, pages 165-178. (2008)

Prices, distinctions: Best paper award DISC 2009.

FAUCONNIER Hugues

Coordinator

Age:55
Position: MC
Email: fauconnier@liafa.jussieu.fr
URL: www.lliafa.jussieu.fr/~hf

LIAFA, Université Paris Diderot
Case 7014
F-75205 Paris cedex 13
Tel: (+33) 1 57 27 92 25

Cursus

PhD (1982) and HDR (2001) from University Denis Diderot

Publications

Number of publications in refereed international journals: 10
(*JACM, TOPLAS, Distributed computing, TCS, Information and Computation, JPDC,TPDS.....)*

Number of publications in refereed international conferences with proceedings: 38
(*PODC, DISC, DSN, ICDCS, SRDS, SPAA, OPODIS, SIROCCO, IPDPS, ICDCN,DCOSS....*)

Selected publications from the past five years
(Authors are given by the alphabetical order of their name)

- C.Delporte-Gallet, H.Fauconnier and R. Guerraoui. **Tight failure detection bounds on atomic object implementations**. *J. ACM*, volume 57 (4). (2010)
- M.K. Aguilera,C.Delporte-Gallet, H.Fauconnier and S. Toueg. **On implementing Omega in systems with weak reliability and synchrony assumptions**. *Distributed Computing*, volume 21 (4), pages 285-314. (2010)
- C.Delporte-Gallet, H.Fauconnier, R. Guerraoui and A. Tielmann. **Fault-Tolerant Consensus in Unknown and Anonymous Networks**. *ICDCS*, pages 368-375). (2009)
- C.Delporte-Gallet, H.Fauconnier, R. Guerraoui and A. Tielmann. **The Disagreement Power of an Adversary**. *DISC(Best Paper)*, pages 8-21). (2009)
- C.Delporte-Gallet, H.Fauconnier and R. Guerraoui. **Sharing is harder than agreeing**. *PODC*, pages 85-94. (2008)

Prices, distinctions: Best paper award DISC 2009.

FRAIGNIAUD Pierre

Age: 48
Position: Directeur de Recherches CNRS
Email: pierre.fraigniaud@liafa.jussieu.fr
URL: www.liafa.jussieu.fr/~pierref

LIAFA, Université Paris Diderot
Case 7014
F-75205 Paris cedex 13
Tel: (+33) 1 57 27 94 00

Cursus

PhD (1990) and Habilitation (1994) from ENS Lyon

Other professional experiences:

P. Fraigniaud is Director of LIAFA (Laboratoire d'Informatique Algorithmique: Fondements et Applications), which includes 60 permanent researchers from University Paris Diderot, CNRS, and INRIA, for a total number of members of roughly 100. Before this charge, he has been vice-Director of LRI at University Paris Sud, as well as Vice-President Research of the Computer Science Dept. of U. Paris Sud. He is currently member of the Editorial Board of Distributed Computing (DC), and Theory of Computing Systems (TOCS), among others. He is the current Program Committee Chair of the 30th Annual ACM Symposium on Principles of Distributed Computing (PODC). Before, he has acted as PC Chair for the 19th Int. Symp. on Distributed Computing (DISC 2005), and for the 13th ACM Symp. on Parallel Algorithms and Architectures (SPAA 2001). He is regular member of PC for conferences such as PODC, SPAA, DISC, ESA, ICALP, WG, MFCS, etc.

Publications

Number of refereed international journal: 58

Number of referreed international conference with proceeding: 88

Selected publications from the past five years
(Authors are given by the alphabetical order of their name)

- Fraigniaud P. and Korman A.. An Optimal Ancestry Scheme and Small Universal Posets. In proc. 42th ACM Symposium on Theory of Computing (STOC), pp 611-620. (2010).
- Fraigniaud P. and G. Giakkoupis. On the searchability of small-world networks with arbitrary underlying sructure. In proc. 42th ACM Symposium on Theory of Computing (STOC), pp 389-398. (2010)
- Baumann H., P. Crescenzi, and Fraigniaud P.. Parsimonious Flooding in Dynamic Graphs. In proc. 28th ACM Symposium on Principles of Distributed Computing (PODC), p. 260-269. (2009)
- Fraigniaud P., and 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), p. 240-249. (2009)
- A. Chaintreau, Fraigniaud P., E. Lebhar. Networks Become Navigable as Nodes Move and Forget. In proc. 35th Int. Colloquium on Automata, Languages and Programming (ICALP), LNCS 5125 , Springer, pp. 133-144. (2008)

Prices, distinctions: Invited Keynote Speaker ACM PODC 2010, ICALP 2010, ICDT 2010, and ESA 2007; Best paper award ACM SPAA 2007.

KORMAN Amos

Age: 38
Position: Chargé de recherches CNRS
Email: amos.korman@liafa.jussieu.fr
URL: www.liafa.jussieu.fr/~pandit

LIAFA, Université Paris Diderot
Case 7014
F-75205 Paris cedex 13
Tel: (+33) 1 57 27 94 00

Cursus

Amos Korman received the PhD degree by Weizmann Institute of Science, Rehovot, Israel on May 8, 2006

Other professional experiences:

Publications

Number of refereed international journal: 17

Number of refereed international conference with proceeding: 26

Selected publications from the past five years

- Emek Y. and Korman A. and Shavitt Y..
  Approximating the Statistics of various Properties in Randomly Weighted Graphs.
  In Proc. 22nd ACM-SIAM Symp. on Discrete Algorithms (SODA), to appear. (2011).
- Fraigniaud P. and Korman A.. An Optimal Ancestry Scheme and Small Universal Posets.
  In proc. 42th ACM Symposium on Theory of Computing (STOC), pp 611-620. (2010).
- Korman A.. Labeling Schemes for Vertex Connectivity.
  In ACM Transactions on Algorithms (TALG), 6(2). (2010).
- Emek Y. and Korman A.. Efficient Threshold Detection in a Distributed Environment.
  In Proc. 29th Symp. on Principles of Distributed Computing (PODC), pp 183-19. (2010).
- Korman A. and Kutten S.. Distributed Verification of Minimum Spanning Trees.
  In Distributed Computing (DC), 20(4), pp 253-266. (2007).
- Korman A.. General Compact Labeling Schemes for Dynamic Trees.
  In Distributed Computing (DC), 20(3), pp 179-193. (2007).

Prices, distinctions:

- Won the 2009 ICDCN best paper award.
- Won the Dean's Prize for Ph.D. Students in Weizmann Institute.
- Won the 2005 DISC best student paper award.
- Graduated at the Hebrew University with exceptional honors: Magna Cum Lauda.

VIENNOT Laurent

Age: 39
Position: DR
Email: laurent.viennot@inria.fr
URL: gang.inria.fr/~viennot/

LIAFA, Université Paris Diderot
Case 7014
F-75205 PARIS cedex 13
Tel: (+33) 1 57 27 94 00

Cursus

Polytechnique, Thèse, Habilitation

Other professional experiences:

Laurent Viennot works on graph and network algorithms. He is a senior research scientist at french national institute on computer science INRIA since 1998. He has also been a part time teacher at Ecole Polytechnique for 7 years. He has also co-founded a startup-company on using peer-to-peer algorithms for photo sharing in 2008. He is now leader of the "GANG" INRIA project-team on network and graph algorithms.

Publications

Number of refereed international journal: 7

Number of referreed international conference with proceeding: 32

Prices, distinctions:

Selected publications from the past five years
(Authors are given by the alphabetical order of their name)

- Gavoille C., Godfroy Q. and Viennot L.. **Multipath Spanners**;. *In* Structural Information and Communication Complexity, 17th International Colloquium (SIROCCO), pages 211–223. (2010)
- Jacquet P. and Viennot L.. **Remote spanners: what to know beyond neighbors** *In* 23rd IEEE International Parallel and Distributed Processing Symposium (IPDPS), pages 1–15. (2009)
- Boufkhad Y., Mathieu F., de Montgolfier F., Perino D. and Viennot L. **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)
- Derbel B., Gavoille C., Peleg D., and Viennot L. **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)
- Lebhar E., Fraigniaud P., and Viennot L. **The inframetric model for the internet.** *In* Proceedings of the 27th IEEE International Conference on Computer Communications (INFOCOM), pages 1–9. (2008)

Partner 2: Rennes

BOURNEZ Olivier

Age: 37
Position: Professor of Computer Science at Ecole Polytechnique
Email: bournez@lix.polytechnique.fr
URL: http://www.lix.polytechnique.fr/~bournez/

LIX, UMR7161
Ecole Polytechnique
Laboratoire d'Informatique
F-91128 Palaiseau Cedex
Tel: (+33) 1 69 33 40 78

Cursus

- **Since 01/09/2008**, **Professor of Computer Science**, Ecole Polytechnique.
  Team Algorthmique et Complexité, lab. LIX (Ecole Polytechnique and CNRS). Director of the lab since 2010.

◦ **01/10/1999 – 31/08/2008**, **Chargé Recherche INRIA** position in Nancy. Habilitation in 2006.

◦ **01/11/1997 – 31/08/1998**, **Scientifique du Contingent**, Laboratoire Verimag, Grenoble.

◦ **01/09/1995 – 14/01/1999**, **Phd in Computer Science** in 1999 in Ecole Normale Supérieure de Lyon. Thèse accessit du Prix de Thse Specif. Distinction par l'AFIT.

◦ **01/09/1992 – 31/08/1996**, **Student of Ecole Normale Supérieure de Lyon**.

Research interests

Computability and Complexity, in particular

- Computability and Complexity with distributed models
- Computability and Complexity with anonymous networks
- Computability and Complexity with analog models
- Algorithmic game theory.

Publications

Number of publications in refereed international journals: 17
(*Journal of Complexity, TCS, PPL, Applied Maths and Computation, Information and Computation, Fundamenta Informaticae, JLC, TOCS, JCSS,*)

Number of publications in referreed international conferences with proceedings: 34
(*ICALP, LCC, UC, MCU, RTA, TAMC, TCS, FOSSACS, HSCC, STACS*)

Selected publications from the past five years
(Authors are given by the alphabetical order of their name)

- G. Aupy and O. Bournez. **On the number of binary-minded individuals required to compute $\sqrt{1/2}$.** *Theoretical Computer Science* to appear. (2010)
- O. Bournez, P. Chassaing, J. Cohen, L. Gerin, and X. Koegler. **On the convergence of population protocols when population goes to infinity**. Applied Mathematics and Computation, 2009. (2009)
- O. Bournez, M. L. Campagnolo, D. S. Graça, and E. Hainry. **Polynomial differential equations compute all real computable functions on computable compact intervals. Journal of Complexity**, 23(3):317 335, June 2007. (2007).
- O. Bournez and M. L. Campagnolo. *New Computational Paradigms. Changing Conceptions of What is Computable*, chapter **A Survey on Continuous Time Computations**, pages 383–423. Springer-Verlag, New York, 2008. (2007).
- D. Barth, O. Bournez, O. Boussaton, and J. Cohen. **Distributed learning of equilibria in a routing game**. *Parallel Processing Letters*, 19:189–204, 2009. (2009).

MOSTEFAOUI Achour

Age: 41
Position: Maître de Conférences
Email: achour.mostefaoui@irisa.fr
URL: www.irisa.fr/asap/

IRISA/IFSIC, Campus de Beaulieu
Avenue du Général Leclerc
35042 Rennes cedex
Tel: (+33) 2 99 84 71 96

Cursus

- **Since 01/09/1996**, **Maîrte de Conférences**.
  Ifsic/Irisa, Université de Rennes.
- ○ **01/09/1994 – 01/09/1996**, **ATER**.
  Ifsic/Irisa, Université de Rennes.
- ○ **01/09/1991 – 01/09/1994**, **PhD in Computer Science**.
  Ifsic/Irisa, Université de Rennes.

Research interests

Distributed algorithms and systems, in particular

- Distributed data structures
- Decidabiliy and efficiency in distributed computing
- Fault-tolerance in distributed systems

Publications

Number of publications in refereed international journals: 32
JACM, Siam J. Comput., Dist. Comp., IEEE TC, IEEE TPDS, IEEE TDSC, JPDC, JCSS, TCS, etc.

Number of publications in referreed international conferences with proceedings: 77
STOC, PODC, DISC, DSN, SPAA, ICDCS, IPDPS, OPODIS, SIROCCO, SRDS, EDCC,etc.

Selected publications from the past five years
(Authors are given by the alphabetical order of their name)

- Achour Mostfaoui, Michel Raynal, Corentin Travers. **Narrowing power vs efficiency in synchronous set agreement: Relationship, algorithms and lower bound**. *Theoreical Computer Science*, vol 411(1), pp. 58-69. (2010)
- Achour Mostfaoui, Sergio Rajsbaum, Michel Raynal, Corentin Travers. **The Combined Power of Conditions and Information on Failures to Solve Asynchronous Set Agreement**. *SIAM J. Comput.*, vol 38(4), pp. 1574-1601. (2008)
- Achour Mostfaoui, Sergio Rajsbaum, Michel Raynal, Corentin Travers. **On the computability power and the robustness of set agreement-oriented failure detector classes**. *Distributed Computing*, vol 21(3), pp. 201-222 (2008)
- Achour Mostfaoui, Sergio Rajsbaum, Michel Raynal. **Synchronous condition-based consensus**. *Distributed Computing*, vol 18(5), pp. 325-343 (2006)
- Roy Friedman, Achour Mostfaoui, Michel Raynal. *Simple and Efficient Oracle-Based Consensus Protocols for Asynchronous Byzantine Systems. IEEE Trans. Dependable Sec. Comput*, vol 2(1), pp. 46-56 (2005)

RAYNAL Michel

Age: 61
Position: Professeur, IUF(Membre Senior)
Email: michel.raynal@irisa.fr
URL: www.irisa.fr/prive/raynal/

IRISA/IFSIC, Campus de Beaulieu
Avenue du Général Leclerc
35042 Rennes cedex
Tel: (+33) 2 99 84 71 96

Cursus

- **Since 01/12/1983**, **Professor**.
  Ifsic/Irisa, Université de Rennes.
- **01/11/1981 – 30/11/1983**, **Professor**.
  Sup Telecom Bretagne, Brest.

Research interests

Distributed algorithms and systems, in particular:

• Decidabiliy and efficiency in distributed computing,
• Fault-tolerance and dependability, • Software transactional memory.

Publications

Number of publications in refereed international journals: 125
(*Journal of the ACM, Algorithmica, SIAM Journal of Computing, Acta Informatica, Distributed Computing, The Communications of the ACM, Information and Computation, Journal of Computer and System Sciences, JPDC, IEEE Trans. on Computers, IEEE Trans. on Software Engineering, IEEE Trans. on Knowledge and Data Engineering, IEEE Trans. on Parallel and Distributed Systems, IEEE Computer, IEEE Software, Journal of Supercomputing, Information Proc. Letters, Parallel Proc. Letters, Theoretical Computer Science, Theory of Computing Systems, Real-Time Systems Journal, The Computer Journal, etc.*)

Prices and distinctions: (*h-index 45, Best paper award: SSS 2009, DISC 2010; Distinguihed paper Europar 2010, Invited Keynote speaker: DISC 2007, AINA 2008*)
Number of publications in referred int'l conferences with proceedings: 258
(*STOC, PODC, DISC, DSN, SPAA, ICDCS, IPDPS, OPODIS, SIROCCO, SRDS, EDCC, ICDCN, etc., etc.*)

Selected publications from the past five years
(Authors are given by the alphabetical order of their name)

- Yehuda Afek, Eli Gafni, Sergio Rajsbaum, Michel Raynal, Corentin Travers. **The k-simultaneous consensus problem**. *Distributed Computing*, vol 22(3), pp. 185-195. (2010)
- Achour Mostfaoui, Michel Raynal, Corentin Travers. **Narrowing power vs efficiency in synchronous set agreement: Relationship, algorithms and lower bound**. *Theoretical Computer Science*, vol 411(1), pp. 58-69. (2010)
- Yoram Moses, Michel Raynal. **Revisiting simultaneous consensus with crash failures**. *Journal of Parallel Distributed Computing*, vol 69(4), pp. 400-409. (2009)
- Achour Mostfaoui, Sergio Rajsbaum, Michel Raynal, Corentin Travers. **On the computability power and the robustness of set agreement-oriented failure detector classes**. *Distributed Computing*, vol 21(3), pp. 201-222. (2008)
- Roy Friedman, Achour Mostfaoui, Michel Raynal. **Simple and Efficient Oracle-Based Consensus Protocols for Asynchronous Byzantine Systems**. *IEEE Transactions Dependable Secure Computing*, vol 2(1), pp. 46-56. (2010)

Partner 3: Bordeaux

GAVOILLE Cyril

Age: 40
Position: PR, University of Bordeaux
Email: gavoille@labri.fr
URL: http://www.labri.fr/perso/gavoille

LaBRI, Université de Bordeaux
351 cours de la Libération
F-33405 Talence cedex
Tel: (+33) 5 40 00 88 12

Cursus

- **Since 2002**, Professor at Bordeaux University, LaBRI
- **1996-2002**, Assistant Professor at Bordeaux University, LaBRI
- **1993-1996**, PhD at Ecole Normale Superieure de Lyon, LIP

Research interests

Distributed graph algorithms, distributed data-structures, distributed distance oracles, quantum information, routing algorithms.

Additional information

Deputy Director of the LaBRI (340 people including 140 faculties) in charge of the scientific affairs in 2009 and 2010, he is currently junior member of the "Institut Universitaire de France" since 2009, a prestigious national status for five years. He is General Chair of PODC 2011 (Principles of Distributed Computing) and was Treasurer in 2010, a top world-class conference in Distributed Computing. He was also co-Chair of ICALP 2010 (Int.'l Coll. on Automata, Languages and Programming), and chairman and co-chair for several workshops (SIROCCO 1999, LOCALITY at DISC 2005, AlgoTel 2003). He has participed in more than 20 international program committee conferences, in particular PODC 2005 and 2006, SPAA 2007, DISC 2007 and 2008, OPODIS 2009, ESA 2011.

Publications

Number of publications in refereed international journals: >40
(including books)

Number of publications in refereed international conferences with proceedings: >90

Prices, distinctions: h-index 27, Best Paper Award for the ACM Int.'l conference SPAA in 2006, IUF Junior Member.

Selected publications from the past five years
(Authors are given by the alphabetical order of their name)

- Fraigniaud P., Gavoille C., Ilcinkas D., Pelc A. **Distributed computing with advice: Information sensitivity of graph coloring**, *Distributed Computing*, 21(6):395-403. (2009)
- Derbel B., Gavoille C., Peleg D., Viennot L. **On the locality of distributed sparse spanner construction**, *27th Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 273-282. (2008)
- Gavoille C., Abraham I., Malkhi D., Nisan N., Thorup M. **Compact name-independent routing with minimum stretch**, *ACM Transactions on Algorithms*, 3(4):A37. (2008)
- Gavoille C., Kosowski A., Markiewicz M. **What can be observed locally? Round-based models for quantum distributed computing**, *23rd International Symposium on Distributed Computing (DISC)*, vol. 5805 of Lecture Notes in Computer Science, pages 243-257. (2009)
- Fraigniaud P., Gavoille C., Kosowski A., Lebhar E., Lotker Z. **Universal Augmentation schemes for network navigability: Overcoming the $\sqrt{n}$-barrier**, *19th Annual ACM Symposium on Parallelism and Architectures (SPAA)*, pages 1-7. (2007)

ILCINKAS David

Age: 30
Position: CNRS junior researcher (CR2)
Email: david.ilcinkas@labri.fr
URL: http://www.labri.fr/perso/ilcinkas

LaBRI, Université de Bordeaux
351 cours de la Libération
F-33405 Talence cedex
Tel: (+33) 5 40 00 69 12

Cursus

- **Since 01/11/2007**, **Junior researcher (CR2)**, CNRS.
  Team Combinatorics and Algorithmics, lab. LaBRI (Université de Bordeaux).
- ◦ **01/09/2006 – 31/08/2007**, **Post-doctoral** position, Université du Québec en Outaouais & University of Ottawa & Carleton University, Canada.
- ◦ **01/10/2003 – 31/08/2006**, **Allocataire-moniteur**, Université Paris-Sud, team GraphComm, LRI, Orsay. **PhD:** in 2006.

Research interests

Distributed graph algorithms, in particular

- Mobile agent computing
- Distributed computing with an oracle / advice

Additional information

Program committee member of several workshops and conferences: SIROCCO 2009, DYNAS 2009 and 2010, AlgoTel 2009 and 2011, DIALM-POMC 2010, DISC 2011.

Publications

Number of publications in refereed international journals: 12
(*ACM TAlg, Algorithmica, DAM, Dist. Comp. [2], Fund. Inf., Inf. & Comp., JCSS, TCS [4]*)

Number of publications in refereed international conferences with proceedings: 22
(*DISC [3], ICALP [3], IWDC, MFCS [2], OPODIS [2], PODC [2], SIROCCO [6], STACS, SWAT, WG*)

Selected publications from the past five years
(Authors are given by the alphabetical order of their name)

- D. Ilcinkas, D. R. Kowalski and A. Pelc. **Fast Radio Broadcasting with Advice**. *Theoretical Computer Science*, volume 411 (14-15), pages 1544-1557. (2010)
- P. Fraigniaud, D. Ilcinkas and A. Pelc. **Communication algorithms with advice**. *Journal of Computer and System Sciences*, volume 76 (3-4), pages 222-232. (2010)
- P. Fraigniaud, C. Gavoille, D. Ilcinkas and A. Pelc. **Distributed computing with advice: information sensitivity of graph coloring**. *Dist. Comp.*, volume 21 (6), pages 395-403. (2009)
- R. Cohen, P. Fraigniaud, D. Ilcinkas, A. Korman and D. Peleg. **Label-Guided Graph Exploration by a Finite Automaton**. *ACM Trans. on Algorithms*, volume 4 (4), article 42. (2008)
- P. Fraigniaud, D. Ilcinkas and A. Pelc. **Tree Exploration with Advice**. *Information and Computation*, volume 206 (11), pages 1276-1287. (2008)

JOHNEN Colette

Age: 49
Position: Full Professor (PR2)
Email: colette.johnen@labri.fr
URL: http://www.labri.fr/∼johnen

LaBRI, Université de Bordeaux
351 cours de la Libération
F-33405 TALENCE cedex
Tel: (+33) 5 40 00 60 47

Cursus

- **Since 01/10/2008**, **Full Professor**, Université de Bordeaux
  Team Combinatorics and Algorithmics, lab. LABRI (CNRS - UMR 5800).

- **01/12/1988 – 31/08/2008**, **Associated Professor** Université de Paris-Sud
  Team Parallelism, lab. LRI (CNRS - UMR 8623).

Research interests

Fault-Tolerant Distributed algorithms

Additional information

PhD Students since 2006 : 3

2004 - 2007 Advisor for 80% of Le Huy Nguyen's PhD thesis (Dr. Pierre Fraigniaud was the co-advisor).
2005 - 2009 Advisor for 20% of Florent Kaisser's PhD thesis (Pr. Véronique Vèque was the co-advisor).
2008 - 2011 Advisor for 80% of Fouzi Mekhaldi's PhD thesis (Pr. Véronique Vèque was the co-advisor).

Publications

Number of publications in refereed international journals : 7
(*TAAS, TCS, Distributed Computing [2], IPL, PPL, JPDC*)

Number of publications in refereed international conferences with proceedings: 34
(*AlgoSensors, ICDCN, ICDCS, IPDPS, EuroPar, SSS, SRDS, OPODIS, DISC, PODC, ...*)

Selected publications from the past five years
(Authors are given by the alphabetical order of their name)

- Colette Johnen and Fouzi Mekhaldi. **Robust Self-stabilizing Construction of Bounded Size Weight-Based Clusters**. *Euro-Par*, LNCS 6271, pages 535-546. (2010)
- Colette Johnen and Lisa Higham. **Fault-Tolerant Implementations of Regular Registers by Safe Registers with Applications to Networks**. *ICDCN*, LNCS 5408, pages 337-348. (2009)
- Kajari Ghosh Dastidar, Ted Herman and Colette Johnen. **Safe peer-to-peer self-downloading**. *TAAS*, vol 3 (4), paper 19. (2008)
- Joffroy Beauquier, Maria Gradinariu and Colette Johnen: **Randomized self-stabilizing and space optimal leader election under arbitrary scheduler on rings**. *Distributed Computing*, vol 20 (1), pages 75-93. (2007)
- Joffroy Beauquier, Colette Johnen, Stéphane Messika. **All k -Bounded Policies Are Equivalent for Self-stabilization**. *SSS* LNCS 4280, pages 82-94. (2006)

KLASING Ralf

Age: 47
Position: CNRS senior researcher (DR2)
Email: klasing@labri.fr
URL: http://www.labri.fr/perso/klasing

LaBRI, Université de Bordeaux
351 cours de la Libération
F-33405 Talence cedex
Tel: (+33) 5 40 00 35 24

Cursus

Ralf Klasing received the PhD degree from the University of Paderborn in 1995. From 1995 to 1997, he was an Assistant Professor at the University of Kiel. From 1997 to 1998, he was a Research Fellow at the University of Warwick. From 1998 to 2000, he was an Assistant Professor at RWTH Aachen. From 2000 to 2002, he was a Lecturer at King's College London. In 2002, he joined the CNRS as a permanent researcher. From 2002 to 2005, he was affiliated to the laboratory I3S in Sophia Antipolis. Currently, he is affiliated to the laboratory LaBRI in Bordeaux. In 2009, he received the HDR degree from the University Bordeaux 1. In 2010, he was promoted to Senior Researcher (DR2 CNRS).

Research interests

Design and Analysis of Algorithms, Communication algorithms in networks, Approximation algorithms for combinatorially hard problems, Algorithmic methods for telecommunication, Distributed algorithms.

Additional information

Head of the *Combinatorics and Algorithms* team of the LaBRI. Member of the Editorial Boards of *Theoretical Computer Science*, *Discrete Applied Mathematics*, *Wireless Networks*, *Networks*, *Journal of Interconnection Networks*, *Parallel Processing Letters*, *Algorithmic Operations Research*, *Fundamenta Informaticae*, *Computing and Informatics*. Member of the Program Committees of *SIROCCO 2006*, *SIROCCO 2009*, *MFCS 2009*, *ADHOC-NOW 2009*, *OPODIS 2009*, *STACS 2010*, *ALGOSENSORS 2010*, *ADHOC-NOW 2010*, *IWOCA 2010*, *SIROCCO 2011*, *FOMC 2011*, *ADHOC-NOW 2011*, *IWOCA 2011*. Member of the Evaluation Committee of the programme *ANR Défis 2009*. Workshops Chair of *ICALP 2010*.

Publications

Number of books: 2 (*Springer Monographs*)
Number of book chapters: 3
Number of publications in refereed international journals: 33 (*SIAM J. Discr. Math., DAM [2], Inf. & Comp. [2], JCSS, TCS [8], IEEE Trans. Par. & Distr. Syst., J. Par. & Distr. Comp., Algorithmica, Networks [4], R.A.I.R.O., ...*)
Number of publications in refereed international conferences with proceedings: 35
(*ICALP, STACS [4], ESA, MFCS, SWAT, WG [2], DISC [2], OPODIS [4], SIROCCO [4], CIAC [2], ISAAC [2], ...*)

Selected publications from the past five years

- L. Gasieniec, R. Klasing, R. Martin, A. Navarra, X. Zhang. **Fast Periodic Graph Exploration with Constant Memory**. *Journal of Computer and System Sciences* 74, No. 5, 808–822. (2008)
- J. Hromkovič, P. Kanarek, R. Klasing, K. Lorys, W. Unger, H. Wagener. **On the Size of Permutation Networks and Consequences for Efficient Simulation of Hypercube Algorithms on Bounded-Degree Networks**. *SIAM Journal on Discrete Mathematics* 23, No. 3, 1612–1645. (2009)
- C. Cooper, D. Ilcinkas, R. Klasing, A. Kosowski. **Derandomizing Random Walks in Undirected Graphs Using Locally Fair Exploration Strategies**. In: Proc. 36th International Colloquium on Automata, Languages and Programming (*ICALP 2009*), *Lecture Notes in Computer Science* 5556, Springer-Verlag, 411–422. (2009)
- C. Cooper, R. Klasing and T. Radzik. **Locating and repairing faults in a network with mobile agents**. *Theoretical Computer Science* 411(14–15):1638–1647. (2010)
- R. Klasing, A. Kosowski and A. Navarra. **Taking Advantage of Symmetries: Gathering of many Asynchronous Oblivious Robots on a Ring**. *Theoretical Computer Science* 411(34–36):3235–3246. (2010)

KOSOWSKI Adrian

Age: 24
Position: INRIA researcher (CR1)
Email: adrian.kosowski@labri.fr
URL: http://www.labri.fr/perso/kosowski

INRIA Bordeaux Sud-Ouest, LaBRI
351 cours de la Libération
F-33405 Talence cedex
Tel: (+33) 5 40 00 69 12

Cursus

- **Since 15/10/2010**, **Researcher (CR1)**, INRIA.
  CEPAGE project, lab. LaBRI (INRIA Bordeaux Sud-Ouest).
- ◦ **01/10/2007 – 31/10/2010**, **Assistant Professor**, Department of Algorithms and System Modeling, Gdańsk University of Technology, Poland.
- ◦ **01/07/2008 – 30/06/2009**, **Post-doctoral** position, Team Combinatorics and Algorithmics, lab. LaBRI (Université de Bordeaux).
- ◦ **01/10/2005 – 30/09/2007**, **Research and teaching assistant**, Department of Algorithms and System Modeling, Gdańsk University of Technology, Poland. **PhD:** in 2007.

Research interests

- **Distributed algorithms**, in particular: distributed graph coloring algorithms, mobile agent computing, self-stabilization and self-organization.
- **Combinatorial optimization algorithms**, in particular: packing and routing paths in graphs, graph coloring, graph exploration, computational geometry.
- **Quantum information** and distributed quantum computing.

Publications

Number of publications in refereed international journals: 19
(*Wireless Networks, Theor. Comp. Sci. [3], Disc. App. Math. [3], Disc. Math. [2], Networks [2], Inf. Process. Lett. [2], Comp. Geom. Theory, and others*)

Number of publications in refereed international conferences with proceedings: 30
(*PODC [2], ICALP, SPAA, MFCS [2], DISC [3], ISAAC, SPIRE, SIROCCO [2], OPODIS [4], Euro-Par, and others*)

Selected publications from the past five years
(Authors are given by the alphabetical order of their name)

- R. Klasing, A. Kosowski, A. Navarra. **Taking advantage of symmetries: Gathering of many asynchronous oblivious robots on a ring**. Theoretical Computer Science, vol 411 (34-36), pp 3235-3246. (2010)
- C. Gavoille, A. Kosowski, M. Markiewicz. **What can be observed locally? Round-based models for quantum distributed computing.** Proceedings of 23rd International Symposium on Distributed Computing, Lecture Notes in Computer Science, vol 5805, pp 243-257. (2009)
- C. Gavoille, R. Klasing, A. Kosowski, L. Kuszner, A. Navarra. **On the complexity of distributed graph coloring with local minimality constraints**. Networks, vol 54 (1), pp 12-19. (2009)
- P. Fraigniaud, C. Gavoille, A. Kosowski, E. Lebhar, Z. Lotker. **Universal augmentation schemes for network navigability**. Theoretical Computer Science, vol 410 (21-23), pp 1970-1981. (2009)
- A. Kosowski. **The maximum edge-disjoint paths problem in complete graphs**. Theoretical Computer Science, vol 399 (1-2), pp 128-140. (2008)

TRAVERS Corentin

Age: 30
Position: Associate professor (Maître de conférence)
Email: travers@labri.fr
URL: www.labri.fr/perso/travers

LaBRI, Université de Bordeaux
351 cours de la Libération
F-33405 Talence cedex
Tel: (+33) 5 40 00 35 29

Cursus

- **Since 01/09/2010**, **Assistant Professor (Maître de conférence)**, ENSEIRB.
  Team Combinatorics and Algorithmics, lab. LaBRI (Université de Bordeaux).
- ◦ **01/10/2009 – 31/08/2010**, **Post-doctoral** position, Université Pierre et Marie Curie.
- ◦ **01/10/2008 – 30/09/2009**, **Post-doctoral** position, The Technion, Israel.
- ◦ **01/01/2008 – 30/09/2008**, **Post-doctoral** position, Universidad Politécnica de Madrid.
- ◦ **01/10/2004 – 31/12/2007**, **Research/Teaching assistant**, Université de Rennes 1, IRISA. **PhD:** in 2007.

Research interests

Theory of distributed Computing, in particular
- Fault-tolerant computing
- Distributed computing with oracles / advice

Additional information

Program committee member of the conference and the workshop: ICDCS 2008 and WRAS 2010.
Won the best paper award of the conference ICDCN 2011 for the paper **Generating Fast Indulgent Algorithms** coauthored with D. Alistarh, S. Gilbert and R. Guerraoui.

Publications

Number of publications in refereed international journals: 10
(*Dist. Comp. [2], Parallel Processing Letters, Trans. Parallel Distrib. Syst., J. Parallel Distrib. Comput., SIAM J. Comput., Inf. Process. Lett., Theor. Comput. Sci. [2], Theory Comput. Syst.*)
Number of publications in refereed international conferences with proceedings: 21
(*COCOON, DISC [3], ICDCN [3], ISAAC, LATIN, OPODIS [3], PODC [2], PRDC [3], SIROCCO [2], SRDS [2]*)

Selected publications from the past five years
(Authors are given by the alphabetical order of their name)

- Y. Afek, E. Gafni, S. Rajsbaum, M. Raynal, C. Travers. **The k-simultaneous Consensus Problem**. *Distributed Computing*, vol 22(3), pp 185-195. (2010)
- A. Mostfaoui, M. Raynal, C. Travers. **Narrowing Power vs Efficiency in Synchronous Set Agreement: Relationship, Algorithms and Lower Bound**. *Theoretical Computer Science*, vol 411 (1), pp 58-69. (2010)
- E. Gafni, A. Mostfaoui, M. Raynal, C. Travers. **From Adaptive Renaming to Set Agreement**. *Theoretical Computer Science*, vol 410(14), pp 1328-1335. (2009)
- A. Mostfaoui, S. Rajsbaum, M. Raynal, C. Travers. **On the Computability Power and the Robustness of Set Agreement-oriented Failure Detector Classes**. *Distributed Computing*, vol 21(3), pp 201-222. (2008)
- A. Mostéfaoui, S. Rajsbaum, M. Raynal, C. Travers. **The Combined Power of Conditions and Information on Failures to Solve Asynchronous Set Agreement**. *SIAM Journal of Computing*, vol 38(4), pp1574-1601. (2008)

7.3 Staff involvment in other contracts

Partner 1: Paris

Part.	Name	PM(to do)	Project name, financing, institution, grant allocated	Project title	Coordinator name	Start and end dates
			ANR VERSO (105 000 E)	SHAMAN	Tixeuil	2008-2012
1	Delporte	2				
1	Fauconnier	7				
1	Fraigniaud	1				
1	Viennot	9,6				
1	Fraigniaud	6	ANR VERSO (104 000 E)	PROSE	Chaintreau	Sept 2009-Aug 2012
1	Korman	3				
1	Viennot	3				

Partner 2: Rennes

Part.	Name	PM	Project name, financing, institution, grant allocated	Project title	Coordinator name	Start and end dates
			ANR VERSO (105 000 E)	SHAMAN	Tixeuil	2008-2012
2	Bournez	19				
2	Kermarrec	5				
2	Mostefaoui	19				
2	Raynal	9				
			ANR ARPÈGE (99 008 E)	STREAMS	Oster	2010-2013
2	Mostefaoui	8				
2	Raynal	8				
			ERC Junior - EC (1 000 000 E)	GOOSPLE	Kermarrec	2008-2013
2	Kermarrec	48				

Partner 3: Bordeaux

Part.	Name	PM	Project name, financing institution, grant allocated	Project title	Coordinator name	Start and end dates
3	Adrian Kosowski	1	Royal Society grant, 12000 pounds	SERENE - SEarch, RENdezvous, and Explore	Leszek Gasieniec	Jan. 1, 2011 - Dec. 31, 2012
3	Ralf Klasing	2	Royal Society grant, 12000 pounds	SERENE - SEarch, RENdezvous, and Explore	Leszek Gasieniec	Jan. 1, 2011 - Dec. 31, 2012
3	David Ilcinkas	1	Royal Society grant, 12000 pounds	SERENE - SEarch, RENdezvous, and Explore	Leszek Gasieniec	Jan. 1, 2011 - Dec. 31, 2012
3	David Ilcinkas	4	FP7 STREP Project, European Commission, 3 150 000 euros	EULER - Experimental Update-Less Evolutive Routing	Dimitri Papadimitriou	Oct. 1, 2010 - Sept. 30, 2013
3	Cyril Gavoille	1	FP7 STREP Project, European Commission, 3 150 000 euros	EULER - Experimental Update-Less Evolutive Routing	Dimitri Papadimitriou	Oct. 1, 2010 - Sept. 30, 2013