



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

Project-Team grand-large

*Calcul parallèle et distribué à grande
échelle*

Saclay - Île-de-France

Theme : Distributed and High Performance Computing

Activity
R *eport*

2009

Table of contents

1. Team	1
2. Overall Objectives	2
3. Scientific Foundations	2
3.1. Large Scale Distributed Systems (LSDS)	2
3.1.1. Computing on Large Scale Global Computing systems	3
3.1.2. Building a Large Scale Distributed System	4
3.1.2.1. The resource discovery engine	4
3.1.2.2. Data storage and movement	5
3.1.2.3. Scheduling in large scale systems	5
3.1.2.4. Fault Tolerant MPI	6
3.2. Volatility and Reliability Processing	7
3.3. Parallel Programming on Peer-to-Peer Platforms (P5)	8
3.3.1. Large Scale Computational Sciences and Engineering	8
3.3.2. Experimentations and Evaluations	8
3.3.3. Languages, Tools and Interface	9
3.4. Methodology for Large Scale Distributed Systems	9
3.4.1. Observation tools	10
3.4.2. Tool for scalability evaluations	10
3.4.3. Real life testbeds: extreme realism	10
3.5. High Performance Scientific Computing	11
3.5.1. Efficient linear algebra algorithms	11
3.5.2. Preconditioning techniques	12
4. Application Domains	12
4.1. Building a Large Scale Distributed System for Computing	12
4.2. Security and Reliability of Network Control Protocols	12
4.3. End-User Tools for Computational Science and Engineering	13
5. Software	13
5.1. APMC-CA	13
5.2. XtremWeb	14
5.3. XtremLab	14
5.4. MPICH-V	15
5.5. YML	16
5.6. The Scientific Programming InterNet (SPIN)	17
5.7. V-DS	17
5.8. PVC: Private Virtual Cluster	18
5.9. OpenWP	18
5.10. FAult Injection Language (FAIL)	19
5.11. Parallel solvers for solving linear systems of equations	19
6. New Results	19
6.1. Non-self-stabilizing and self-stabilizing gathering in networks of mobile agents—the notion of speed	19
6.2. Making Population Protocols Self-stabilizing	20
6.3. Self-stabilizing distributed control	20
6.4. Extensions of Self-stabilization	22
6.5. Byzantine Resilience in autonomous networks	22
6.6. Self-Stabilizing Control Infrastructure for HPC	23
6.7. Large Scale Peer to Peer Performance Evaluations	23
6.7.1. Large Scale Grid Computing	24
6.7.2. High Performance Cluster Computing	24

6.7.3. Large Scale Power aware Computing	24
6.8. High Performance Linear Algebra on the Grid	24
6.9. Emulation of Volatile Systems	25
6.10. Exascale Systems	25
6.11. High performance scientific computing	26
6.11.1. Communication avoiding algorithms for LU and QR factorizations	26
6.11.2. Preconditioning techniques	27
7. Other Grants and Activities	27
7.1.1. Activities starting in 2009	27
7.1.2. Other activities	28
8. Dissemination	28
8.1.1. Research Administration	28
8.1.2. Book/Journal edition	28
8.1.3. Conference Organisation	29
8.1.4. Editorial Committee membership	29
8.1.5. Steering Committee membership	29
8.1.6. Program Committee membership	29
9. Bibliography	30

1. Team

Research Scientist

Franck Cappello [Team Leader, Research Director, HdR]

Laura Grigori [Junior Researcher CR1, HdR]

Faculty Member

Joffroy Beauquier [Professor at Paris-Sud University, HdR]

Thomas Héroult [Assistant Professor at Paris Sud University. Delegated in the INRIA project/team]

Serge Petitot [Professor at University of Science and Technology of Lille, HdR]

Sylvain Peyronnet [Assistant Professor at Paris Sud University]

Brigitte Rozoy [Professor at Paris-Sud University, HdR]

Sébastien Tixeuil [Professor at Paris 6 University, HdR]

Technical Staff

Roberto Podesta [INRIA Expert Engineer, end September 2009]

Matthieu Cargnelli [INRIA Expert Engineer May 2009]

Victor Iriarte [INRIA Expert Engineer, end October 2009]

Vincent Néri [CNRS Study Engineer]

PhD Student

Ali Asim [MESR Grant LRI]

Fatiha Bouabache [MESR Grant (LRI)]

Camille Coti [INRIA Grant, end November 2009]

Pawan Kumar [INRIA Grant]

Thomas Largillier [MESR Grant (LRI)]

Paul Malécot [INRIA Grant]

Ala Rezmerita [MESR Grant (LRI), end September 2009]

Amina Guermouche [MESR Grant (LRI)]

Alexandre Borghi [MESR Grant (LRI)]

Simplice Donfack [INRIA Grant]

Ye Zhang [Since 2006]

Ling Shang [French Ministry of Foreign Affairs, since 2007]

Amal Khabou [MESR Grant (LRI), since October 2009]

Antoine Baldacci [CIFRE IFP, since November 2009]

Post-Doctoral Fellow

Alok Gupta [PostDoc, end February 2009]

Ala Rezmerita [MESR Grant (LRI), since October 2009]

Federico Stivoli [PostDoc, since November 2009]

Guy Antoine Atenekeng Kahou [PostDoc, since January 2009]

Elisabeth Brunet [PostDoc, since April 2009]

Ke Wang [PostDoc, since April 2009]

Francois Lesueur [MESR Grant (LRI), since October 2009]

Administrative Assistant

Katia Evrat [Administrative assistant]

2. Overall Objectives

2.1. Grand-Large General Objectives

Grand-Large is a research project investigating the issues raised by High Performance Computing (HPC) on Large Scale Distributed Systems (LSDS), where users execute HPC applications on a shared infrastructure and where resources are subject to failure, possibly heterogeneous, geographically distributed and administratively independent. More specifically, we consider large scale distributed computing mainly, Desktop Grids, Grids, and large scale parallel computers. Our research focuses on the design, development, proof and experiments of programming environments, middleware and scientific algorithms and libraries for HPC applications. Fundamentally, we address the issues related to HPC on LSDS, gathering several methodological tools that raise themselves scientific issues: theoretical models and exploration tools (simulators, emulators and real size experimental systems).

Our approach ranges from concepts to experiments, the projects aims at:

1. models and fault-tolerant algorithms, self-stabilizing systems and wireless networks.
2. studying experimentally, and formally, the fundamental mechanisms of LSDS for high performance computing;
3. designing, implementing, validating and testing real software, libraries, middleware and platforms;
4. defining, evaluating and experimenting approaches for programming applications on these platforms.

Compared to other European and French projects, we gather skills in 1) large scale systems formal design and validation of algorithms and protocols for distributed systems and 2) programming, evaluation, analysis and definition of programming languages and environments for parallel architectures and distributed systems.

This project pursues short and long term researches aiming at having scientific and industrial impacts. Research topics include:

1. the design of middleware for LSDS (XtremWeb and PVC)
2. large scale data movements on LSDS (BitDew)
3. fault tolerant MPI for LSDS, fault tolerant protocol verification (MPICH-V)
4. algorithms, programming and evaluation of scientific applications LSDS;
5. tools and languages for large scale computing on LSDS (OpenWP, YML).
6. Exploration systems and platforms for LSDS (Grid'5000, XtremLab, DSL-Lab, SimBOINC, FAIL, V-DS)

These researches should have some applications in the domain of Desktop Grids, Grids and large scale parallel computers.

As a longer term objective, we put special efforts on the design, implementation and use of Exploration Tools for improving the methodology associated with the research in LSDS. For example we had the responsibility of the Grid eXplorer project founded by the French ministry of research and we were deeply involved in the Grid5000 project (as project Director) and in the ALADDIN initiative (project scientific director).

3. Scientific Foundations

3.1. Large Scale Distributed Systems (LSDS)

What makes a fundamental difference between recent Global Computing systems (Seti@home), Grid (EGEE, TeraGrid) and former works on distributed systems is the large scale of these systems. This characteristic becomes also true for large scale parallel computers gathering tens of thousands of CPU cores. The notion of

Large Scale is linked to a set of features that has to be taken into account in these systems. An example is the system dynamicity caused by node volatility: in Internet Computing Platforms (also called Desktop Grids), a non predictable number of nodes may leave the system at any time. Some recent results also report a very low MTTI (Mean Time To Interrupt) in top level supercomputers gathering 100,000+ CPU cores. Another example of characteristics is the complete lack of control of nodes connectivity. In Desktop Grid, we cannot assume that external administrator is able to intervene in the network setting of the nodes, especially their connection to Internet via NAT and Firewalls. This means that we have to deal with the in place infrastructure in terms of performance, heterogeneity, dynamicity and connectivity. These characteristics, associated with the requirement of scalability, establish a new research context in distributed systems. The Grand-Large project aims at investigating theoretically as well as experimentally the fundamental mechanisms of LSDS, especially for the high performance computing applications.

3.1.1. Computing on Large Scale Global Computing systems

Large scale parallel and distributed systems are mainly used in the context of Internet Computing. As a consequence, until Sept. 2007, Grand-Large has focused mainly on Desktop Grids. Desktop Grids are developed for computing (SETI@home, Folding@home, Decryphon, etc.), file exchanges (Napster, Kazaa, eDonkey, Gnutella, etc.), networking experiments (PlanetLab, Porivo) and communications such as instant messaging and phone over IP (Jabber, Skype). In the High Performance Computing domain, LSDS have emerged while the community was considering clustering and hierarchical designs as good performance-cost tradeoffs. Nowadays, Internet Computing systems are still very popular (the BOINC platform is used to run over 40 Internet Computing projects and XtremWeb is used in production in three countries) and still raise important research issues.

Desktop Grid systems essentially extend the notion of computing beyond the frontier of administration domains. The very first paper discussing this type of systems [122] presented the Worm programs and several key ideas that are currently investigated in autonomous computing (self replication, migration, distributed coordination, etc.). LSDS inherit the principle of aggregating inexpensive, often already in place, resources, from past research in cycle stealing/resource sharing. Due to its high attractiveness, cycle stealing has been studied in many research projects like Condor [106], Glunix [97] and Mosix [71], to cite a few. A first approach to cross administration domains was proposed by Web Computing projects such as Jet [110], Charlotte [72], Javeline [90], Bayanihan [118], SuperWeb [68], ParaWeb [78] and PopCorn [82]. These projects have emerged with Java, taking benefit of the virtual machine properties: high portability across heterogeneous hardware and OS, large diffusion of virtual machine in Web browsers and a strong security model associated with bytecode execution. Performance and functionality limitations are some of the fundamental motivations of the second generation of Global Computing systems like COSM [81], BOINC [70] and XtremWeb [93]. The second generation of Global Computing systems appeared in the form of generic middleware which allow scientists and programmers to design and set up their own distributed computing project. As a result, we have seen the emergence of large communities of volunteers and projects. Currently, Global Computing systems are among the largest distributed systems in the world. In the mean time, several studies succeeded to understand and enhance the performance of these systems, by characterizing the system resources in term of volatility and heterogeneity and by studying new scheduling heuristics to support new classes of applications: data-intensive, long running application with checkpoint, workflow, soft-real time etc... However, despite these recent progresses, one can note that Global Computing systems are not yet part of high performance solution, commonly used by scientists. Recent researches to fulfill the requirements of Desktop Grids for high demanding users aim at redesigning Desktop Grid middleware by essentially turning a set of volatile nodes into a virtual cluster and allowing the deployment of regular HPC utilities (batch schedulers, parallel communication libraries, checkpoint services, etc...) on top of this virtual cluster. The new generation would permit a better integration in the environment of the scientists such as computational Grids, and consequently, would broaden the usage of Desktop Grid.

The high performance potential of LSDS platforms has also raised a significant interest in the industry. Companies like United Devices [127], Platform [112], Grid systems [133] and Datasynapse [132] propose LSDS middleware for Desktop Grids (also known as PC Grid systems). Performance demanding users are

also interested by these platforms, considering their cost-performance ratio which is even lower than the one of clusters. Thus, several Desktop Grid platforms are daily used in production in large companies in the domains of pharmacology, petroleum, aerospace, etc.

Desktop Grids share with Grid a common objective: to extend the size and accessibility of a computing infrastructure beyond the limit of a single administration domain. In [94], the authors present the similarities and differences between Grid and Global Computing systems. Two important distinguishing parameters are the user community (professional or not) and the resource ownership (who own the resources and who is using them). From the system architecture perspective, we consider two main differences: the system scale and the lack of control of the participating resources. These two aspects have many consequences, at least on the architecture of system components, the deployment methods, programming models, security (trust) and more generally on the theoretical properties achievable by the system.

Beside Desktop Grids and Grids, large scale parallel computers with tens of thousands (and even hundreds of thousands) of CPU cores are emerging with scalability issues similar to the one of Internet Computing systems: fault tolerance at large scale, large scale data movements, tools and languages. Grand-Large is gradually considering the application of selected research results, in the domain of large scale parallel computers, in particular for the fault tolerance and language topics.

3.1.2. Building a Large Scale Distributed System

This set of studies considers the XtremWeb project as the basis for research, development and experimentation. This LSDS middleware is already operational. This set gathers 4 studies aiming at improving the mechanisms and enlarging the functionalities of LSDS dedicated to computing. The first study considers the architecture of the resource discovery engine which, in principle, is close to an indexing system. The second study concerns the storage and movements of data between the participants of a LSDS. In the third study, we address the issue of scheduling in LSDS in the context of multiple users and applications. Finally the last study seeks to improve the performance and reduce the resource cost of the MPICH-V fault tolerant MPI for desktop grids.

3.1.2.1. The resource discovery engine

A multi-users/multi-applications LSDS for computing would be in principle very close to a P2P file sharing system such as Napster [119], Gnutella [119] and Kazaa [105], except that the shared resource is the CPUs instead of files. The scale and lack of control are common features of the two kinds of systems. Thus, it is likely that solutions sharing fundamental mechanisms will be adopted, such as lower level communication protocols, resource publishing, resource discovery and distributed coordination. As an example, recent P2P projects have proposed distributed indexing systems like CAN [115], CHORD [124], PASTRY [117] and TAPESTRY [131] that could be used for resource discovery in a LSDS dedicated to computing.

The resource discovery engine is composed of a publishing system and a discovery engine, which allow a client of the system to discover the participating nodes offering some desired services. Currently, there is as much resource discovery architectures as LSDS and P2P systems. The architecture of a resource discovery engine is derived from some expected features such as speed of research, speed of reconfiguration, volatility tolerance, anonymity, limited use of the network, matching between the topologies of the underlying network and the virtual overlay network.

This study focuses on the first objective: to build a highly reliable and stable overlay network supporting the higher level services. The overlay network must be robust enough to survive unexpected behaviors (like malicious behaviors) or failures of the underlying network. Unfortunately it is well known that under specific assumptions, a system cannot solve even simple tasks with malicious participants. So, we focus the study on designing overlay algorithms for transient failures. A transient failure accepts any kind of behavior from the system, for a limited time. When failures stop, the system will eventually provide its normal service again.

A traditional way to cope with transient failures are self-stabilizing systems [92]. Existing self-stabilizing algorithms use an underlying network that is not compatible with LSDS. They assume that processors know their list of neighbors, which does not fit the P2P requirements. Our work proposes a new model for designing self-stabilizing algorithms without making this assumption, then we design, prove and evaluate overlay networks self-stabilizing algorithms in this model.

3.1.2.2. Data storage and movement

Application data movements and storage are major issues of LSDS since a large class of computing applications requires the access of large data sets as input parameters, intermediary results or output results.

In this scenario of data-centric applications, the existing Desktop Grid Systems face a scalability issue. One should expect that more computing resources also provides more network bandwidth and storage capacity. On the contrary, Desktop Grids Systems like BOINC or XtremWeb rely on a centralized data service architecture. For instance, data distribution with BOINC relies on multiple HTTP servers and tasks are described as a list of files locations, which can be a potential bottleneck when scheduling tasks sharing large input files. Recent developments in content distribution such as collaborative file distribution (BitTorrent [91], Slurpie [121], Digital Fountain [80]) and P2P file sharing applications (EDonkey/Overnet [66], Kazaa [105]), has proven to be both effective, viable and *self-scalable*. The key idea, often featured as *parallel download* in the P2P applications, is to divide the file in a list of chunks. Immediately after a peer downloads a chunk from another peer, the peer serves the block for the other peers in the network, thus behaving itself as a server. Collaborative Content Distribution is a very active research topic and several promising strategies [113] such as the use of network coding [98], are proposed to improve the performance of the system. We will evaluate the efficiency of collaborative data distribution mechanism according to its impact on the overall performance of a parallel application when scheduled on Desktop Grid.

Currently there is no LSDS system dedicated to computing that allows the persistent storage of data in the participating nodes. Several LSDS systems dedicated to data storage are emerging such as OCEAN Store [102] and Ocean [87]. Storing large data sets on volatile nodes requires replication techniques. In CAN and Freenet, the documents are stored in a single piece. In OceanStore, Fastrack and eDonkey, the participants store segments of documents. This allows segment replications and the simultaneous transfer of several documents segments. In the CGP2P project, a storage system called US has been proposed. It relies on the notion of blocs (well known in hard disc drivers). Redundancy techniques complement the mechanisms and provide raid like properties for fault tolerance. Moreover, we believe that the basic blocks for building a system dedicated to data management in Desktop Grids can be found in P2P systems. We will investigate new paradigm for LSDS data management based on P2P components such as Bittorrent for data distribution and DHT for efficient and scalable data index/search, and design a new middleware, BitDew, featuring data transfer reliability, volatility tolerance, automatic data replication, multi-protocol capability, data affinity, and tuple-space like semantic.

3.1.2.3. Scheduling in large scale systems

Scheduling is one of the system fundamental mechanisms. Several studies have been conducted in the context of Grid mostly considering bag of tasks, parameter sweep or workflow applications [85], [83]. Recently some researches consider scheduling and migrating MPI applications on Grid [123]. Other related researches concern scheduling for cycle stealing environments [116]. Some of these studies consider not only the dynamic CPU workload but also the network occupation and performance as basis for scheduling decisions. They often refer to NWS which is a fundamental component for discovering the dynamic parameters of a Grid. There are very few researches in the context of LSDS and no existing practical ways to measure the workload dynamics of each component of the system (NWS is not scalable).

While LSDS systems consist of volatile and heterogeneous computing resources, it is unknown exactly how volatile and heterogeneous these computing resources are. While there have been previous characterization studies on Internet-wide computing resources, these studies do not take into account causes of volatility and only report coarse aggregate statistics, such as the mean time to failure of resources. Yet, detailed resource characterization is essential for the simulation and modelling of LSDS systems in a research area where many results are obtained via simulation.

We propose to design, implement, and deploy a resource monitoring project called XtremLab via the BOINC software system. The goal of XtremLab will be to monitor the availability of a large fraction of the LSDS participants in an effort to paint a more detailed picture of the Internet computing landscape.

Based, on these availability traces, we will design new simulator of Global Computing system (SIMBOINC) and investigate advanced scheduling heuristics for bag of task applications with large input data sets,

application with soft real time constraints, workflow and data flow applications, long running applications which needs periodic checkpoints etc...

3.1.2.4. Fault Tolerant MPI

MPICH-V is a research effort with theoretical studies, experimental evaluations and pragmatic implementations aiming to provide a MPI implementation based on MPICH [108], featuring multiple fault tolerant protocols.

There is a long history of research in fault tolerance for distributed systems. We can distinguish the automatic/transparent approach from the manual/user controlled approach. The first approach relies either on coordinated checkpointing (global snapshot) or uncoordinated checkpointing associated with message logging. A well known algorithm for the first approach has been proposed by Chandy and Lamport [86]. This algorithm requires restarting all processes even if only one process crashes. So it is believed not to scale well. Several strategies have been proposed for message logging: optimistic [128], pessimistic [69], causal [129]. Several optimizations have been studied for the three strategies. The general context of our study is high performance computing on large platforms. One of the most used programming environments for such platforms is MPI.

Within the MPICH-V project, we have developed and published several original fault tolerant protocols for MPI: MPICH-V1 [75], MPICH-V2 [76], MPICH-Vcausal, MPICH-Vcl [77], MPICH-Pcl. The two first protocols rely on uncoordinated checkpointing associated with either remote pessimistic message logging or sender based pessimistic message logging. We have demonstrated that MPICH-V2 outperforms MPICH-V1. MPICH-Vcl implements a coordinated checkpoint strategy (Chandy-Lamport) removing the need of message logging. MPICH-V2 and Vcl are concurrent protocols for large clusters. We have compared them considering a new parameter for evaluating the merits of fault tolerant protocols: the impact of the fault frequency on the performance. We have demonstrated that the stress of the checkpoint server is the fundamental source of performance differences between the two techniques. MPICH-Vcausal implements a causal message logging protocols, removing the need for waiting acknowledgement in contrary to MPICH-V2. MPICH-Pcl is a blocking implementation of the Vcl protocol. Under the considered experimental conditions, message logging becomes more relevant than coordinated checkpoint when the fault frequency reaches 1 fault every 4 hours, for a cluster of 100 nodes sharing a single checkpoint server, considering a data set of 1 GB on each node and a 100 Mb/s network.

Multiple important events arose from this research topic. A new open source implementation of the MPI-2 standard was born during the evolution of the MPICH-V project, namely OpenMPI. OpenMPI is the result of the alliance of many MPI projects in the USA, and we are working to port our fault tolerance algorithms both into OpenMPI and MPICH.

Grids becoming more popular and accessible than ever, parallel applications developers now consider them as possible targets for computing demanding applications. MPI being the de-facto standard for the programming of parallel applications, many projects of MPI for the Grid appeared these last years. We contribute to this new way of using MPI through a European Project in which we intend to grid-enable OpenMPI and provide new fault-tolerance approaches fitted for the grid.

When introducing Fault-Tolerance in MPI libraries, one of the most neglected component is the runtime environment. Indeed, the traditional approach consists in restarting the whole application and runtime environment in case of failure. A more efficient approach could be to implement a fault-tolerant runtime environment, capable of coping with failures at its level, thus avoiding the restart of this part of the application. The benefits would be a quicker restart time, and a better control of the application. However, in order to build a fault-tolerant runtime environment for MPI, new topologies, more connected, and more stable, must be integrated in the runtime environment.

For traditional parallel machines of large scale (like large scale clusters), we also continue our investigation of the various fault tolerance protocols, by designing, implementing and evaluating new protocols in the MPICH-V project.

3.2. Volatility and Reliability Processing

In a global computing application, users voluntarily lend the machines, during the period they don't use them. When they want to reuse the machines, it is essential to give them back immediately. We assume that there is no time for saving the state of the computation (for example because the user is shooting down is machine). Because the computer may not be available again, it is necessary to organize checkpoints. When the owner takes control of his machine, one must be able to continue the computation on another computer from a checkpoint as near as possible from the interrupted state.

The problems raised by this way of managing computations are numerous and difficult. They can be put into two categories: synchronization and repartition problems.

- Synchronization problems (example). Assume that the machine that is supposed to continue the computation is fixed and has a recent checkpoint. It would be easy to consider that this local checkpoint is a component of a global checkpoint and to simply rerun the computation. But on one hand the scalability and on the other hand the frequency of disconnections make the use of a global checkpoint totally unrealistic. Then the checkpoints have to be local and the problem of synchronizing the recovery machine with the application is raised.
- Repartition problems (example). As it is also unrealistic to wait for the computer to be available again before rerunning the interrupted application, one has to design a virtual machine organization, where a single virtual machine is implemented as several real ones. With too few real machines for a virtual one, one can produce starvation; with too many, the efficiency is not optimal. The good solution is certainly in a dynamic organization.

These types of problems are not new ([95]). They have been studied deeply and many algorithmic solutions and implementations are available. What is new here and makes these old solutions not usable is scalability. Any solution involving centralization is impossible to use in practice. Previous works validated on former networks can not be reused.

3.2.1. Reliability Processing

We voluntarily presented in a separate section the volatility problem because of its specificity both with respect to type of failures and to frequency of failures. But in a general manner, as any distributed system, a global computing system has to resist to a large set of failures, from crash failures to Byzantine failures, that are related to incorrect software or even malicious actions (unfortunately, this hypothesis has to be considered as shown by DECRYPTHON project or the use of erroneous clients in SETI@HOME project), with in between, transient failures such as loss of message duplication. On the other hand, failures related accidental or malicious memory corruptions have to be considered because they are directly related to the very nature of the Internet. Traditionally, two approaches (masking and non-masking) have been used to deal with reliability problems. A masking solution hides the failures to the user, while a non-masking one may let the user notice that failures occur. Here again, there exists a large literature on the subject (cf. [107], [125], [92] for surveys). Masking techniques, generally based on consensus, are not scalable because they systematically use generalized broadcasting. The self-stabilizing approach (a non-masking solution) is well adapted (specifically its time adaptive version, cf. [104], [103], [73], [74], [96]) for three main reasons:

1. Low overhead when stabilized. Once the system is stabilized, the overhead for maintaining correction is low because it only involves communications between neighbours.
2. Good adaptivity to the reliability level. Except when considering a system that is continuously under attacks, self-stabilization provides very satisfying solutions. The fact that during the stabilization phase, the correctness of the system is not necessarily satisfied is not a problem for many kinds of applications.
3. Lack of global administration of the system. A peer to peer system does not admit a centralized administrator that would be recognized by all components. A human intervention is thus not feasible and the system has to recover by itself from the failures of one or several components, that is precisely the feature of self-stabilizing systems.

We propose:

1. To study the reliability problems arising from a global computing system, and to design self-stabilizing solutions, with a special care for the overhead.
2. For problem that can be solved despite continuously unreliable environment (such as information retrieval in a network), to propose solutions that minimize the overhead in space and time resulting from the failures when they involve few components of the system.
3. For most critical modules, to study the possibility to use consensus based methods.
4. To build an adequate model for dealing with the trade-off between reliability and cost.

3.3. Parallel Programming on Peer-to-Peer Platforms (P5)

Several scientific applications, traditionally computed on classical parallel supercomputers, may now be adapted for geographically distributed heterogeneous resources. Large scale P2P systems are alternative computing facilities to solve grand challenge applications.

Peer-to-Peer computing paradigm for large scale scientific and engineering applications is emerging as a new potential solution for end-user scientists and engineers. We have to experiment and to evaluate such programming to be able to propose the larger possible virtualization of the underlying complexity for the end-user.

3.3.1. Large Scale Computational Sciences and Engineering

Parallel and distributed scientific application developments and resource managements in these environments are a new and complex undertaking. In scientific computation, the validity of calculations, the numerical stability, the choices of methods and software are depending of properties of each peer and its software and hardware environments; which are known only at run time and are non-deterministic. The research to obtain acceptable frameworks, methodologies, languages and tools to allow end-users to solve accurately their applications in this context is capital for the future of this programming paradigm.

GRID scientific and engineering computing exists already since more than a decade. Since the last few years, the scale of the problem sizes and the global complexity of the applications increase rapidly [126]. The scientific simulation approach is now general in many scientific domains, in addition to theoretical and experimental aspects, often link to more classic methods. Several applications would be computed on world-spread networks of heterogeneous computers using some web-based Application Server Provider (ASP) dedicated to targeted scientific domains. New very strategic domains, such as Nanotechnologies, Climatology or Life Sciences, are in the forefront of these applications. The development in this very important domain and the leadership in many scientific domains will depend in a close future to the ability to experiment very large scale simulation on adequate systems [120], [101]. The P2P scientific programming is a potential solution, which is based on existing computers and networks. The present scientific applications on such systems are only concerning problems which are mainly data independents: i.e. each peer does not communicate with the others.

P2P programming has to develop parallel programming paradigms which allow more complex dependencies between computing resources. This challenge is an important goal to be able to solve large scientific applications. The results would also be extrapolated toward future petascale heterogeneous hierarchically designed supercomputers.

3.3.2. Experimentations and Evaluations

We have followed two tracks. First, we did experiments on large P2P platforms in order to obtain a realistic evaluation of the performance we can expect. Second, we have set some hypothesis on peers, networks, and scheduling in order to have theoretical evaluations of the potential performance. Then, we have chosen a classical linear algebra method well-adapted to large granularity parallelism and asynchronous scheduling: the block Gauss-Jordan method to invert dense very large matrices. We have also chosen the calculation of one matrix polynomial, which generates computation schemes similar to many linear algebra iterative

methods, well-adapted for very large sparse matrices. Thus, we were able to theoretically evaluate the potential throughput with respect to several parameters such as the matrix size and the multicast network speed.

Since the beginning of the evaluations, we experimented with those parallel methods on a few dozen peer XtremWeb P2P Platforms. We continue these experiments on larger platforms in order to compare these results to the theoretical ones. Then, we would be able to extrapolate and obtain potential performance for some scientific applications.

Recently, we also experimented several Krylov based method, such as the Lanczos and GMRES methods on several grids, such as a French-Japanese grid using hundred of PC in France and 4 clusters at the University of Tsukuba. We also experimented on GRID5000 the same methods. We currently use several middleware such as Xtremweb, OmniRPC and Condor. We also begin some experimentations on the Tsubame supercomputer in collaboration with the TITech (Tokyo Institute of Technologies) in order to compare our grid approaches and the High performance one on an hybrid supercomputer.

Experimentations and evaluation for several linear algebra methods for large matrices on P2P systems will always be developed all along the Grand Large project, to be able to confront the different results to the reality of the existing platforms.

As a challenge, we would like, in several months, to efficiently invert a dense matrix of size one million using a several thousand peer platform. We are already inverting very large dense matrices on Grid5000 but more efficient scheduler and a larger number of processors are required to this challenge.

Beyond the experimentations and the evaluations, we propose the basis of a methodology to efficiently program such platforms, which allow us to define languages, tools and interface for the end-user.

3.3.3. Languages, Tools and Interface

The underlying complexity of the Large Scale P2P programming has to be mainly virtualized for the end-user. We have to propose an interface between the end-user and the middleware which may extract the end-user expertise or propose an on-the-shelf general solution. Targeted applications concern very large scientific problems which have to be developed using component technologies and up-to-dated software technologies.

We introduced the YML framework and language which allows to describe dependencies between components. We introduced different classes of components, depending of the level of abstraction, which are associated with divers parts of the framework. A component catalogue is managed by an administrator and/or the end-users. Another catalogue is managed with respect to the experimental platform and the middleware criteria. A front-end part is completely independent of any middleware or testbed, and a back-end part is developed for each targeted middleware/platform couple. A YML scheduler is adapted for each of the targeted systems.

The YML framework and language propose a solution to develop scientific applications to P2P and GRID platform. An end-user can directly develop programs using this framework. Nevertheless, many end-users would prefer avoid programming at the component and dependency graph level. Then, an interface has to be proposed soon, using the YML framework. This interface may be dedicated to a special scientific domain to be able to focus on the end-user vocabulary and P2P programming knowledge. We plan to develop such version based on the YML framework and language. The first targeted scientific domain will be very large linear algebra for dense or sparse matrices.

3.4. Methodology for Large Scale Distributed Systems

Research in the context of LSDS involves understanding large scale phenomena from the theoretical point of view up to the experimental one under real life conditions.

One key aspects of the impact of large scale on LSDS is the emergence of phenomena which are not coordinated, intended or expected. These phenomena are the results of the combination of static and dynamic features of each component of LSDS: nodes (hardware, OS, workload, volatility), network (topology, congestion, fault), applications (algorithm, parameters, errors), users (behavior, number, friendly/aggressive).

Validating current and next generation of distributed systems targeting large-scale infrastructures is a complex task. Several methodologies are possible. However, experimental evaluations on real testbeds are unavoidable in the life-cycle of a distributed middleware prototype. In particular, performing such real experiments in a rigorous way requires to benchmark developed prototypes at larger and larger scales. Fulfilling this requirement is mandatory in order to fully observe and understand the behaviors of distributed systems. Such evaluations are indeed mandatory to validate (or not!) proposed models of these distributed systems, as well as to elaborate new models. Therefore, to enable an experimentally-driven approach for the design of next generation of large scale distributed systems, developing appropriate evaluation tools is an open challenge.

Fundamental aspects of LSDS as well as the development of middleware platforms are already existing in Grand-Large. Grand-Large aims at gathering several complementary techniques to study the impact of large scale in LSDS: observation tools, simulation, emulation and experimentation on real platforms.

3.4.1. Observation tools

Observation tools are mandatory to understand and extract the main influencing characteristics of a distributed system, especially at large scale. Observation tools produce data helping the design of many key mechanisms in a distributed system: fault tolerance, scheduling, etc. We pursue the objective of developing and deploying a large scale observation tool (XtremLab) capturing the behavior of thousands of nodes participating to popular Desktop Grid projects. The collected data will be stored, analyzed and used as reference in a simulator (SIMBOINC).

3.4.2. Tool for scalability evaluations

Several Grid and P2P systems simulators have been developed by other teams: SimGrid [84], GridSim [79], Briks [67]. All these simulators considers relatively small scale Grids. They have not been designed to scale and simulate 10 K to 100 K nodes. Other simulators have been designed for large multi-agents systems such as Swarm [109] but many of them considers synchronous systems where the system evolution is guided by phases. In the P2P field, ad hoc many simulators have been developed, mainly for routing in DHT. Emulation is another tool for experimenting systems and networks with a higher degree of realism. Compared to simulation, emulation can be used to study systems or networks 1 or 2 orders of magnitude smaller in terms of number of components. However, emulation runs the actual OS/middleware/applications on actual platform. Compared to real testbed, emulation considers conducting the experiments on a fully controlled platform where all static and dynamic parameters can be controlled and managed precisely. Another advantage of emulation over real testbed is the capacity to reproduce experimental conditions. Several implementations/configurations of the system components can be compared fairly by evaluating them under the similar static and dynamic conditions. Grand-Large is leading one of the largest Emulator project in Europe called Grid explorer (French funding). This project has built and used a 1K CPUs cluster as hardware platform and gathers 24 experiments of 80 researchers belonging to 13 different laboratories. Experiments concerned developing the emulator itself and use of the emulator to explore LSDS issues. In term of emulation tool, the main outcome of Grid explorer is the V-DS system, using virtualization techniques to fold a virtual distributed system 50 times larger than the actual execution platform. V-DS aims at discovering, understanding and managing implicit uncoordinated large scale phenomena. Grid Explorer is still in use within the Grid'5000 platform and serves the community of 400 users 7 days a week and 24h a day.

3.4.3. Real life testbeds: extreme realism

The study of actual performance and connectivity mechanisms of Desktop Grids needs some particular testbed where actual middleware and applications can be run under real scale and real life conditions. Grand-Large is developing DSL-Lab, an experimental platform distributed on 50 sites (actual home of the participants) and using the actual DSL network as the connection between the nodes. Running experiments over DSL-Lab put the piece of software to study under extremely realistic conditions in terms of connectivity (NAT, Firewalls), performance (node and network), performance symmetry (DSL Network is not symmetric), etc.

To investigate real distributed system at large scale (Grids, Desktop Grids, P2P systems), under real life conditions, only a real platform (featuring several thousands of nodes), running the actual distributed system can provide enough details to clearly understand the performance and technical limits of a piece of software. Grand-Large members are strongly involved (as Project Director) in the French Grid5000 project which intends to deploy an experimental Grid testbed for computer scientists. This testbed features about 4000 CPUs gathering the resources of about 9 clusters geographically distributed over France. The clusters will be connected by a high speed network (Renater 10G). Grand-Large is the leading team in Grid5000, chairing the steering committee. As the Principal Investigator of the project, Grand-Large has taken some strong design decisions that nowadays give a real added value of Grid5000 compared to all other existing Grids: reconfiguration and isolation. From these two features, Grid5000 provides the capability to reproduce experimental conditions and thus experimental results, which is the cornerstone of any scientific instrument.

3.5. High Performance Scientific Computing

This research is in the area of high performance scientific computing, and in particular in parallel matrix algorithms. This is a subject of crucial importance for numerical simulations as well as other scientific and industrial applications, in which linear algebra problems arise frequently. The modern numerical simulations coupled with ever growing and more powerful computational platforms have been a major driving force behind a progress in numerous areas as different as fundamental science, technical/technological applications, life sciences.

The main focus of this research is on the design of efficient, portable linear algebra algorithms, such that solving a large set of linear equations or computing eigenvalues and eigenvectors. The characteristics of the matrices commonly encountered in this situations can vary significantly, as are the computational platforms used for the calculations.

Nonetheless two common trends are easily discernible. First, the problems to solve are larger and larger, since the numerical simulations are using higher resolution. Second, the architecture of today's supercomputers is getting very complex, and so the developed algorithms need to be adapted to these new architectures.

A number of methods and solvers exist for solving linear systems. They can be divided into three classes: direct, iterative or semi-iterative. Direct methods (LU factorization for solving linear systems and QR factorization for solving least squares problems) are often preferred because of their robustness. The methods differ significantly depending on whether the matrices are dense (all nonzero entries) or sparse (very few nonzero entries, common in matrices arising from physical modelling). Iterative methods as Krylov subspace iterations are less robust, but they are widely used because of their limited memory requirements and good scalability properties on sparse matrices. Preconditioners are used to accelerate the convergence of iterative methods. Semi-iterative methods such as subdomain methods are hybrid direct/iterative methods which can be good tradeoffs.

3.5.1. Efficient linear algebra algorithms

For the last several years, we have worked on a novel approach to dense and sparse linear algebra algorithms, which aims at minimizing the communication, in terms of both its volume and a number of transferred messages. This research is motivated by technological trends showing an increasing communication cost. Its main goal is to reformulate and redesign linear algebra algorithms so that they are optimal in an amount of the communication they perform, while retaining the numerical stability. The work here involves both theoretical investigation and practical coding on diverse computational platforms.

The theoretical investigation focuses on identifying lower bounds on communication for different operations in linear algebra, where communication refers to data movement between processors in the parallel case, and to data movement between different levels of memory hierarchy in the sequential case. The lower bounds are used to study the existing algorithms, understand their communication bottlenecks, and design new algorithms that attain them. The results obtained to date concern the LU, QR and rank revealing QR factorizations of dense matrices.

This research focuses on the design of linear algebra algorithms that minimize the cost of communication. Communication costs include both latency and bandwidth, whether between processors on a parallel computer or between memory hierarchy levels on a sequential machine. The stability of the new algorithms represents an important part of this work.

3.5.2. Preconditioning techniques

Solving a sparse linear system of equations is the most time consuming operation at the heart of many scientific applications, and therefore it has received a lot of attention over the years. While direct methods are robust, they are often prohibitive because of their time and memory requirements. Iterative methods are widely used because of their limited memory requirements, but they need an efficient preconditioner to accelerate their convergence. In this direction of research we focus on preconditioning techniques for solving large sparse systems.

4. Application Domains

4.1. Building a Large Scale Distributed System for Computing

The main application domain of the Large Scale Distributed System developed in Grand-Large is high performance computing. The two main programming models associated with our platform (RPC and MPI) allow to program a large variety of distributed/parallel algorithms following computational paradigms like bag of tasks, parameter sweep, workflow, dataflow, master worker, recursive exploration with RPC, and SPMD with MPI. The RPC programming model can be used to execute concurrently different applications codes, the same application code with different parameters and library function codes. In all these cases, there is no need to change the code. The code must only be compiled for the target execution environment. LSDS are particularly useful for users having large computational needs. They could typically be used in Research and Development departments of Pharmacology, Aerospace, Automotive, Electronics, Petroleum, Energy, Meteorology industries. LSDS can also be used for other purposes than CPU intensive applications. Other resources of the connected PCs can be used like their memory, disc space and networking capacities. A Large Scale Distributed System like XtremWeb can typically be used to harness and coordinated the usage of these resources. In that case XtremWeb deploys on Workers services dedicated to provide and manage a disc space and the network connection. The storage service can be used for large scale distributed fault tolerant storage and distributed storage of very large files. The networking service can be used for server tests in real life conditions (workers deployed on Internet are coordinated to stress a web server) and for networking infrastructure tests in real like conditions (workers of known characteristics are coordinated to stress the network infrastructure between them).

4.2. Security and Reliability of Network Control Protocols

The main application domain for self-stabilizing and secure algorithms is LSDS where correct behaviours must be recovered within finite time. Typically, in a LSDS (such as a high performance computing system), a protocol is used to control the system, submit requests, retrieve results, and ensure that calculus is carried out accordingly to its specification. Yet, since the scale of the system is large, it is likely that nodes fail while the application is executing. While nodes that actually perform the calculus can fail unpredictably, a self-stabilizing and secure control protocol ensures that a user submitting a request will obtain the corresponding result within (presumably small) finite time. Examples of LSDS where self-stabilizing and secure algorithms are used, include global computing platforms, or peer to peer file sharing systems. Another application domain is routing protocols, which are used to carry out information between nodes that are not directly connected. Routing should be understood here in its most general acceptance, e.g. at the network level (Internet routing) or at the application level (on virtual topologies that are built on top of regular topologies in peer to peer systems). Since the topology (actual or virtual) evolves quickly through time, self-stabilization ensures that the routing protocol eventually provides accurate information. However, for the protocol to be useful, it is necessary that

it provides extra guarantees either on the stabilization time (to recover quickly from failures) or on the routing time of messages sent when many faults occur. Finally, additional applications can be found in distributed systems that are composed of many autonomous agents that are able to communicate only to a limited set of nodes (due to geographical or power consumption constraints), and whose environment is evolving rapidly. Examples of such systems are wireless sensor networks (that are typically large of 10000+ nodes), mobile autonomous robots, etc. It is completely unrealistic to use centralized control on such networks because they are intrinsically distributed; still strong coordination is required to provide efficient use of resources (bandwidth, battery, etc).

4.3. End-User Tools for Computational Science and Engineering

Another Grand Large application domain is Linear Algebra, which is often required to solve Large Scale Computational Science and Engineering applications. Two main approaches are proposed. First, we have to experiment and evaluate several classical stable numerical methods. Second, we have to propose tools to help end-users to develop such methods.

In addition to the classical supercomputing and the GRID computing, the large scale P2P approach proposes new computing facilities for computational scientists and engineers. Thus, it exists many applications which would use such computing facilities for long period of time. During a first period, many applications will be based on large simulations rather than classical implicit numerical methods, which are more difficult to adapt for such large problems and new programming paradigm as they generated linear algebra problems. Then, implicit method would be developed to have more accurate solutions.

Simulations and large implicit methods always have to compute linear algebra routines. So, they were our first targeted numerical methods (we also remark that the powerful worldwide computing facilities are still rated using a linear algebra benchmark <http://www.top500.org>). We especially focused on divide-and-conquer and block-based matrix methods to solve dense problems. We have also studied Krylov subspace methods (Lanczos, Arnoldi) and hybrid methods to solve sparse matrix problems. As these applications are utilized for many applications, it is possible to extrapolate the results to different scientific domains.

Many smart tools have to be developed to help the end-user to program such environments, using up-to-date component technologies and languages. At the actual present stage of maturity of this programming paradigm for scientific applications, the main goal is to experiment on large platforms, to evaluate and extrapolate performance, and to propose tools for the end-users; with respect to many parameters and under some specify hypothesis concerning scheduling strategies and multicast speeds [100]. We have to always replace the end-user at the center of this scientific programming. Then, we have to propose a framework to program P2P architectures which completely virtualizes the P2P middleware and the heterogeneous hardware. Our approach is based, on the one hand, on component programming and coordination languages, and on the other hand, to the development of an ASP, which may be dedicated to a targeted scientific domain. The YML framework provides a solution to the first point since it offers the YvetteML workflow language in order to orchestrate components. This is a very intuitive programming approach and it favors the re-usability of optimized and bug-free components. The abstraction of the underlying P2P middleware is also ensured by YML by means of its back-end mechanism. The end-user of YML can submit a computing task to any kind of peer connected to Internet as long as YML has a back-end in charge of the middleware which is running on this peer. Currently, YML has two back-ends for the XtremWeb and OmniRPC middleware. Another one for Condor will be soon available. The second point concerns the integration of SPIN to YML in order to get a complete programming tool which covers all the needs of the client in order to run applications (based on linear algebra methods) over the Internet. Finally, the conclusion of our work would be a P2P scientific programming methodology based on experimentations and evaluation on an actual P2P development environment.

5. Software

5.1. APMC-CA

Participant: Sylvain Peyronnet [correspondant].

APMC-CA (Cell Assisted Approximate Probabilistic Model Checker) is a new version of the APMC model checker specially dedicated for the Cell processor architecture. Using the specific features of a Cell architecture, it achieves better performances than APMC 3.0 (the previous version of the software) running on powerful standard workstations.

5.2. XtremWeb

Participant: Gilles Fedak [correspondant].

XtremWeb is an open source middleware, generalizing global computing platforms for a multi-user and multi-parallel programming context. XtremWeb relies on the notion of services to deploy a Desktop Grid based on a 3 tiers architecture. This architecture gathers three main services: Clients, Coordinators and Workers. Clients submit requests to the coordinator which uses the worker resources to execute the corresponding tasks. XtremWeb decouples access to data from tasks, which allows integration with network file service.

Coordinator sub-services provide resource discovery, service construction, service instantiation and data repository for parameters and results. A major concern is fault tolerance. XtremWeb relies on passive replication and message logging to tolerate Clients mobility, Coordinator transient and definitive crashes and Worker volatility. An extension of XtremWeb called RPC-V features failure-tolerance of the coordinator. The Client service provides a Java API which unifies the interactions between the applications and the Coordinator. Three client applications are available: the Java API that can be used in any Java applications, a command line (shell like) interface and a web interface allowing users to easily submit requests, consult status of their tasks and retrieve results. A second major issue is the security. The origins of the threats are the applications, the infrastructure, the data (parameters and results) and the participating nodes. Currently XtremWeb provides user right management, application sandboxing and communication encryption. XtremWeb provides a RPC interface for bag of tasks, parameter sweep, master worker and workflow applications.

XtremWeb has been tested extensively harnessing a thousand of Workers and computing a million of tasks. XtremWeb is deployed in several sites: LIFL, LIP, ID-IMAG, LRI, LAL (physics laboratory of Orsay), IBBMC (biology laboratory of Orsay), Université de Guadeloupe, IFP (Institut Français du Pétrole), EADS, CEA, University of Wisconsin Madison, University of Tsukuba (Japan), AIST (Australia), UCSD (USA). A Dutch company (AlmereGrid) maintain a Desktop Grid based on XtremWeb in the town of Almere. University of Paris Sud and University of California, San Diego have used XtremWeb to gather hosts availability traces. To our knowledge, several applications has been ported to XtremWeb: Aires belonging to the HEP Auger project (LAL), a protein conformation predictor using a molecular dynamic simulator (IBBMC), CFD code (IFP), PHYLIP Phylogenetic Package and Hand Writing recognition Application (Univ Tunisie), Application Based Iterative Methods (Hohai University, China). Several XtremWeb extension have been developed by research teams: XtremWeb-CH allows direct P2P communication between workers, YvetteML/XtremWeb allows XtremWeb to be used with a workflow frontend.

XtremWeb is available at <http://www.xtremweb.net/>

5.3. XtremLab

Participants: Gilles Fedak [correspondant], Paul Malécot.

Since the late 1990's, DG systems, such as SETI@Home, have been the largest and most powerful distributed computing systems in the world, offering an abundance of computing power at a fraction of the cost of dedicated, custom-built supercomputers. Many applications from a wide range of scientific domains –including computational biology, climate prediction, particle physics, and astronomy– have utilized the computing power offered by DG systems. DG systems have allowed these applications to execute at a huge scale, often resulting in major scientific discoveries that would otherwise had not been possible.

The computing resources that power DG systems are shared with the owners of the machines. Because the resources are volunteered, utmost care is taken to ensure that the DG tasks do not obstruct the activities of each machine's owner; a DG task is suspended or terminated whenever the machine is in use by another person. As a result, DG resources are volatile in the sense that any number of factors can cause the task of a DG application to not complete. These factors include mouse or keyboard activity, the execution of other user applications, machine reboots, or hardware failures. Moreover, DG resources are heterogeneous in the sense that they differ in operating systems, CPU speeds, network bandwidth, memory and disk sizes. Consequently, the design of systems and applications that utilize these systems is challenging.

The long-term overall goal of XtremLab is to create a testbed for networking and distributed computing research. This testbed will allow for computing experiments at unprecedented scale (i.e., thousands of nodes or more) and accuracy (i.e., nodes that are at the "ends" of the Internet).

Currently, the short-term goal of XtremLab is to determine a more detailed picture of the Internet computing landscape by measuring the network and CPU availability of many machines. While DG systems consist of volatile and heterogeneous computing resources, it is unknown exactly how volatile and heterogeneous these computing resources are. Previous characterization studies on Internet-wide computing resources have not taken into account causes of volatility such as mouse and keyboard activity, other user applications, and machine reboots. Moreover, these studies often only report coarse aggregate statistics, such as the mean time to failure of resources. Yet, detailed resource characterization is essential for determining the usefulness of DG systems for various types of applications. Also this characterization is a prerequisite for the simulation and modelling of DG systems in a research area where many results are obtained via simulation, which allow for controlled and repeatable experimentation.

For example, one direct application of the measurements is to create a better BOINC CPU scheduler, which is the software component responsible for distributing tasks of the application to BOINC clients. We plan to use our measurements to run trace-driven simulations of the BOINC CPU scheduler in effort to identify ways it can be improved, and for testing new CPU schedulers before they are widely deployed.

We conduct availability measurements by submitting real compute-bound tasks to the BOINC DG system. These tasks are executed only when the host is idle, as determined by the user's preferences and controlled the BOINC client. These tasks continuously perform computation and periodically record their computation rates to file. These files are collected and assembled to create a continuous time series of CPU availability for each participating host. Utmost care will be taken to ensure the privacy of participants. Our simple, active trace method allows us to measure exactly what actual compute power a real, compute-bound application would be able to exploit. Compared to other passive measurement techniques, our method is not as susceptible to OS idiosyncrasies (e.g. with process scheduling) and takes into account keyboard and mouse activity, and host load, all of which directly impact application execution.

The XtremLab project is available at <http://xtremlab.lri.fr>

5.4. MPICH-V

Participant: Thomas Hérault [correspondant].

Currently, MPICH-V proposes 7 protocols: MPICH-V1, MPICH-V2, MPICH-Vcl, MPICH-Pcl and 3 algorithms for MPICH-Vcausal. MPICH-V1 implements an original fault tolerant protocol specifically developed for Desktop Grids relying on uncoordinated checkpoint and remote pessimistic message logging. It uses reliable nodes called Channel Memories to store all in transit messages. MPICH-V2 is designed for homogeneous networks like clusters where the number of reliable component assumed by MPICH-V1 is too high. It reduces the fault tolerance overhead and increases the tolerance to node volatility. This is achieved by implementing a new protocol splitting the message logging into message payload logging and event logging. These two elements are stored separately on the sender node for the message payload and on a reliable event logger for the message events. The third protocol, called MPICH-Vcl, is derived from the Chandy-Lamport global snapshot algorithm. It implements coordinated checkpoint without message logging. This protocol exhibits

less overhead than MPICH-V2 for clusters with low fault frequencies. MPICH-Pcl is a blocking implementation of Chandy-Lamport algorithm. It consists in exchanging messages for emptying every communication channel before checkpointing all processes. MPICH-Vcausal concludes the set of message logging protocols, implementing a causal logging. It provides less synchrony than the pessimistic logging protocols, allowing messages to influence the system before the sender can be sure that non deterministic events are logged, to the cost of appending some information to every communication. This sum of information may increase with the time, and different causal protocols, with different cut techniques, have been studied with the MPICH-V project.

The protocols developed during the first phase of the MPICH-V project are now being integrated into the two main open-source distributions of MPI, namely MPICH2 and OpenMPI. During this integration, we focus on keeping the best performances (i.e. introducing the smallest changes in the library communication driver). Eventually, the fault-tolerance properties of these two distributions should be provided by the Grand-Large project.

In addition to fault tolerant properties, MPICH-V:

1. provides a full runtime environment detecting and re-launching MPI processes in case of faults;
2. works on high performance networks such as Myrinet, Infiniband, etc (the performances are still divided by two);
3. allows the migration of a full MPI execution from one cluster to another, even if they are using different high performance networks.

The software, papers and presentations are available at <http://mpich-v.lri.fr/>

5.5. YML

Participants: Serge Petiton [correspondant], Nahid Emad.

Scientific end-users face difficulties to program P2P large scale applications using low level languages and middleware. We provide a high level language and a set of tools designed to develop and execute large coarse grain applications on peer-to-peer systems. Thus, we introduced, developed and experimented the YML for parallel programming on P2P architectures. This work was done in collaboration with the PRiSM laboratory (team of Nahid Emad).

The main contribution of YML is its high level language for scientific end-users to develop parallel programs for P2P platforms. This language integrates two different aspects. The first aspect is a component description language. The second aspect allows to link components together. A coordination language called YvetteML can express graphs of components which represent applications for peer-to-peer systems.

Moreover, we designed a framework to take advantage of the YML language. It is based on two component catalogues and an YML engine. The first one concerns end-user's components and the second one is related to middleware criteria. This separation enhances portability of applications and permits real time optimizations. Currently we provide support for the XtremWeb Peer-to-Peer middleware and the OmniRPC grid system. The support for Condor is currently under development and a beta-release will be delivered soon (in this release, we plan to propagate semantic data from the end-users to the middleware). The next development of YML concerns the implementation of a multi-backend scheduler. Therefore, YML will be able to schedule at runtime computing tasks to any global computing platform using any of the targeted middleware.

We experimented YML with basic linear algebra methods on a XtremWeb P2P platform deployed between France and Japan. Recently, we have implemented complex iterative restarted Krylov methods, such as Lanczos-Bisection, GMRES and MERAM methods, using YML with the OmniRPC back-end. The experiments are performed either on the Grid5000 testbed or on a Network of Workstations deployed between Lille, Versailles and Tsukuba in Japan. Demos was proposed on these testbeds from conferences in USA. We recently finished evaluations of the overhead generated using YML, without smart schedulers and with extrapolations due to the lack of smart scheduling strategies inside targeted middleware.

The software is available at <http://yml.prism.uvsq.fr/>

5.6. The Scientific Programming InterNet (SPIN)

Participant: Serge Petiton [correspondant].

SPIN (Scientific Programming on the InterNet), is a scalable, integrated and interactive set of tools for scientific computations on distributed and heterogeneous environments. These tools create a collaborative environment allowing the access to remote resources.

The goal of SPIN is to provide the following advantages: Platform independence, Flexible parameterization, Incremental capacity growth, Portability and interoperability, and Web integration. The need to develop a tool such as SPIN was recognized by the GRID community of the researchers in scientific domains, such as linear algebra. Since the P2P arrives as a new programming paradigm, the end-users need to have such tools. It becomes a real need for the scientific community to make possible the development of scientific applications assembling basic components hiding the architecture and the middleware. Another use of SPIN consists in allowing to build an application from predefined components ("building blocks") existing in the system or developed by the developer. The SPIN users community can collaborate in order to make more and more predefined components available to be shared via the Internet in order to develop new more specialized components or new applications combining existing and new components thanks to the SPIN user interface.

SPIN was launched at ASCI CNRS lab in 1998 and is now developed in collaboration with the University of Versailles, PRiSM lab. SPIN is currently under adaptation to incorporate YML, cf. above. Nevertheless, we study another solution based on the Linear Algebra Kernel (LAKE), developed by the Nahid Emad team at the University of Versailles, which would be an alternative to SPIN as a component oriented integration with YML.

5.7. V-DS

Participant: Franck Cappello [correspondant].

This project started officially in September 2004, under the name V-Grid. V-DS stands for Virtualization environment for large-scale Distributed Systems. It is a virtualization software for large scale distributed system emulation. This software allows folding a distributed systems 100 or 1000 times larger than the experimental testbed. V-DS virtualizes distributed systems nodes on PC clusters, providing every virtual node its proper and confined operating system and execution environment. Thus compared to large scale distributed system simulators or emulators (like MicroGrid), V-DS virtualizes and schedules a full software environment for every distributed system node. V-DS research concerns emulation realism and performance.

A first work concerns the definition and implementation of metrics and methodologies to compare the merits of distributed system virtualization tools. Since there is no previous work in this domain, it is important to define what and how to measure in order to qualify a virtualization system relatively to realism and performance. We defined a set of metrics and methodologies in order to evaluate and compared virtualization tools for sequential system. For example a key parameter for the realism is the event timing: in the emulated environment, events should occur with a time consistent with a real environment. An example of key parameter for the performance is the linearity. The performance degradation for every virtual machine should evolve linearly with the increase of the number of virtual machines. We conducted a large set of experiments, comparing several virtualization tools including Vserver, VMware, User Mode Linux, Xen, etc. The result demonstrates that none of them provides both enough isolation and performance. As a consequence, we are currently studying approaches to cope with these limits.

We have made a virtual platform on the GDX cluster with the Vserver virtualization tool. On this platform, we have launched more than 20K virtual machines (VM) with a folding of 100 (100 VM on each physical machine). However, some recent experiments have shown that a too high folding factor may cause a too long execution time because of some problems like swapping. Currently, we are conducting experiments on another platform based on the virtualization tool named Xen which has been strongly improved since 2 years. We expect to get better result with Xen than with Vserver. Recently, we have been using the V-DS version based on Xen to evaluate at large scales three P2P middleware [114].

This software is available at <http://v-ds.lri.fr/>

5.8. PVC: Private Virtual Cluster

Participant: Franck Cappello [correspondant].

Current complexity of Grid technologies, the lack of security of Peer-to-Peer systems and the rigidity of VPN technologies make sharing resources belonging to different institutions still technically difficult.

We propose a new approach called "Instant Grid" (IG), which combines various Grid, P2P and VPN approaches, allowing simple deployment of applications over different administration domains. Three main requirements should be fulfilled to make Instant Grids realistic: simple networking configuration (Firewall and NAT), no degradation of resource security, no need to re-implement existing distributed applications.

Private Virtual Cluster, is a low-level middle-ware that meets Instant Grid requirements. PVC turns dynamically a set of resources belonging to different administration domains into a virtual cluster where existing cluster runtime environments and applications can be run. The major objective of PVC is to establish direct connections between distributed peers. To connect firewall protected nodes in the current implementation, we have integrated three techniques: UPnP, TCP/UDP Hole Punching and a novel technique Traversing-TCP.

One of the major application of PVC is the third generation desktop Grid middleware. Unlike BOINC and XtremWeb (which belong to the second generation of desktop Grid middleware), PVC allows the users to build their Desktop Grid environment and run their favorite batch scheduler, distributed file system, resource monitoring and parallel programming library and runtime software. PVC ensures the connectivity layer and provide a virtual IP network where the user can install and run existing cluster software.

By offering only the connectivity layer, PVC allows to deploy P2P systems with specific applications, like file sharing, distributed computing, distributed storage and archive, video broadcasting, etc.

5.9. OpenWP

Participant: Franck Cappello [correspondant].

Distributed applications can be programmed on the Grid using workflow languages, object oriented approaches (Proactive, IBIS, etc), RPC programming environments (Grid-RPC, DIET), component based environments (generally based on Corba) and parallel programming libraries like MPI.

For high performance computing applications, most of the existing codes are programmed in C, Fortran and Java. These codes have 100,000 to millions of lines. Programmers are not inclined to rewrite them in a "non standard" programming language, like UPC, CoArray Fortran or Global Array. Thus environments like MPI and OpenMPI remain popular even if they require hybrid approaches for programming hierarchical computing infrastructures like cluster of multi-processors equipped with multi-core processors.

Programming applications on the Grid add a novel level in the hierarchy by clustering the cluster of multi-processors. The programmer will face strong difficulties in adapting or programming a new application for these runtime infrastructures featuring a deep hierarchy. Directive based parallel and distributed computing is appealing to reduce the programming difficulty by allowing incremental parallelization and distribution. The programmer add directives on a sequential or parallel code and may check for every inserted directive its correction and performance improvement.

We believe that directive based parallel and distributed computing may play a significant role in the next years for programming High performance parallel computers and Grids. We have started the development of OpenWP. OpenWP is a directive based programming environment and runtime allowing expressing workflows to be executed on Grids. OpenWP is compliant with OpenMP and can be used in conjunction with OpenMP or hybrid parallel programs using MPI + OpenMP.

The OpenWP environment consists in a source to source compiler and a runtime. The OpenWP parser, interprets the user directives and extracts functional blocks from the code. These blocks are inserted in a library distributed on all computing nodes. In the original program, the functional blocks are replaced by RPC calls and calls to synchronization. During the execution, the main program launches non blocking RPC calls to functions on remote nodes and synchronize the execution of remote functions based on the synchronization directives inserted by the programmer in the main code. Compared to OpenMP, OpenWP does not consider a shared memory programming approach. Instead, the source to source compiler insert data movements calls in the main code. Since the data set can be large in Grid application, the OpenWP runtime organize the storage of data sets in a distributed way. Moreover, the parameters and results of RPC calls are passed by reference, using a DHT. Thus, during the execution, parameter and result references are stored in the DHT along with the current position of the datasets. When a remote function is called, the DHT is consulted to obtain the position of the parameter data sets in the system. When a remote function terminates its execution, it stores the result data sets and store a reference to the data set in the DHT.

We are evaluating OpenWP from an industrial application (Amibe), used by the European aerospace company EADS. Amibe is the mesher module of jCAE¹. Amibe generates a mesh from a CAD geometry in three steps. It first creates edges between every patch of the CAD (mesh in one dimension), then generates a surface mesh for every unfolded patch (mesh in two dimensions) and finally adds the third dimension to the mesh by projecting the 2D mesh into the original CAD surfaces. The first and third operation cannot be distributed. However the second step can easily be distributed following a master/worker approach, transferring the meshId results to every computing node and launching the distributed execution of the patches.

5.10. Fault Injection Language (FAIL)

Participant: Sébastien Tixeuil [correspondant].

FAIL (FAult Injection Language) is a new project that was started in 2004. The goal of this project is to provide a controllable fault injection layer in existing distributed applications (for clusters and grids). A new language was designed to implement expressive fault patterns, and a preliminary implementation of the distributed fault injector based on this language was developed.

5.11. Parallel solvers for solving linear systems of equations

Participant: Laura Grigori.

In the last several years, there has been significant research effort in the development of fully parallel direct solvers for computing the solution of large unsymmetric sparse linear systems of equations. In this context, we have designed and implemented a parallel symbolic factorization algorithm, which is suitable for general sparse unsymmetric matrices. The symbolic factorization is one of the steps that is sequential and represents a memory bottleneck. The code is intended to be used with very large matrices when because of the memory usage, the sequential algorithm is not suitable. This code is available in the SuperLU_DIST, a widely used software, developed at UC Berkeley and LBNL by Professor James W. Demmel and Dr. Xiaoye S. Li. The algorithm is presented in [99]. The SuperLU_DIST is available at <http://crd.lbl.gov/~xiaoye/SuperLU/>.

We continue the development in implementing the numerical factorization phase using the communication avoiding ideas developed this year.

6. New Results

6.1. Non-self-stabilizing and self-stabilizing gathering in networks of mobile agents—the notion of speed

Participants: Joffroy Beauquier, Janna Burman, Julien Clément, Shay Kutten.

¹project page: <http://jcae.sourceforge.net>

In the population protocol model, each agent is represented by a finite state machine. Agents are anonymous and supposed to move in an asynchronous way. When two agents come into range of each other (“meet”), they can exchange information. One of the vast variety of motivating examples to the population protocols model is ZebraNet. ZebraNet is a habitat monitoring application where sensors are attached to zebras and collect biometric data (e.g. heart rate, body temperature) and information about their behavior and migration patterns (via GPS). The population protocol model is, in some sense, related to cloud computing and to networks characterized by asynchrony, large scale, the possibility of failures, in the agents as well as in the communications, with the constraint that each agent is resource limited.

In order to extend the computation power and efficiency of the population protocol model, various extensions were suggested. Our contribution is an extension of the population protocol model that introduces the notion of “speed”, in order to capture the fact that the mobile agents move at different speeds and/or have different communication ranges and/or move according to different patterns and/or visit different places with different frequencies. Intuitively, fast agents which carry sensors with big communication ranges communicate with other agents more frequently than other agents do. This notion is formalized by allocating a cover time, cv , to each mobile agent v . cv is the minimum number of events in the whole system that occur before agent v meets every other agent at least once. As a fundamental example, we have considered the basic problem of gathering information that is distributed among anonymous mobile agents and where the number of agents is unknown. Each mobile agent owns a sensed input value and the goal is to communicate the values (as a multi-set, one value per mobile agent) to a fixed non-mobile base station (BS), with no duplicates or losses.

Gathering is a building block for many monitoring applications in networks of mobile agents. For example, a solution to this problem can solve a transaction commit/abort task in MANETs, if the input values of agents are votes (and the number of agents is known to BS). Moreover, the gathering problem can be viewed as a formulation of the routing problem in Disruption Tolerant Networks.

We gave different solutions to the gathering in the model of mobile agents with speed and we proved that one of them is optimal.

6.2. Making Population Protocols Self-stabilizing

Participants: Joffroy Beauquier, Janna Burman, Shay Kutten, Brigitte Rozoy.

As stated in the previous paragraph, the application domains of the population protocol model are asynchronous large scale networks, in which failures are possible and must be taken into account. This work concerns failures and namely the technique of self-stabilization for tolerating them.

Developing self-stabilizing solutions (and proving them) is considered to be more challenging and complicated than developing classical solutions, where a proper initialization of the variables can be assumed. This remark holds for a large variety of models and hence, to ease the task of the developers, some automatic techniques have been proposed to transform programs into self-stabilizing ones.

We have proposed such a transformer for algorithms in the population protocol model introduced for dealing with resource-limited mobile agents. The model we consider is a variation of the original one in that there is a non mobile agent, the base station, and that the communication characteristics (e.g. moving speed, communication radius) of the agents are considered through the notion of cover time.

The automatic transformer takes as an input an algorithm solving a static problem and outputs a self-stabilizing solution for the same problem. To the best of our knowledge, it is the first time that such a transformer for self-stabilization is presented in the framework of population protocols. We prove that the transformer we propose is correct and we make the complexity analysis of the stabilization time.

6.3. Self-stabilizing distributed control

Participant: Sébastien Tixeuil.

We presented in [52] a 50 pages book chapter devoted to self-stabilizing algorithms in the Handbook of Theory of Computing and Algorithms, published by Taylor and Francis.

We generalized [16] the classic dining philosophers problem to separate the conflict and communication neighbors of each process. Communication neighbors may directly exchange information while conflict neighbors compete for the access to the exclusive critical section of code. This generalization is motivated by a number of practical problems in distributed systems including problems in wireless sensor networks. We presented a self-stabilizing deterministic algorithm — KDP that solves a restricted version of the generalized problem where the conflict set for each process is limited to its k -hop neighborhood. Our algorithm is terminating. We formally proved KDP correct and evaluated its performance. We then extended KDP to handle fully generalized problem. We further extended it to handle a similarly generalized drinking philosophers problem. We described how KDP can be implemented in wireless sensor networks and demonstrated that this implementation does not jeopardize its correctness or termination properties.

We quantified [20] the amount of “practical information (*i.e.* views obtained from the neighbors, colors attributed to the nodes and links) to obtain “theoretical information (*i.e.* the local topology of the network up to distance k) in anonymous networks. In more details, we show that a coloring at distance $2k + 1$ is necessary and sufficient to obtain the local topology at distance k that includes outgoing links. This bound drops to $2k$ when outgoing links are not needed. A second contribution deals with color bootstrapping (from which local topology can be obtained using the aforementioned mechanisms). On the negative side, we showed that (i) with a distributed daemon, it is impossible to achieve deterministic color bootstrap, even if the whole network topology can be instantaneously obtained, and (ii) with a central daemon, it is impossible to achieve distance m when instantaneous topology knowledge is limited to $m - 1$. On the positive side, we showed that (i) under the k -central daemon, deterministic self-stabilizing bootstrap of colors up to distance k is possible provided that k -local topology can be instantaneously obtained, and (ii) under the distributed daemon, probabilistic self-stabilizing bootstrap is possible for any range.

In [36], our focus was to lower the communication complexity of self-stabilizing protocols *below* the need of checking every neighbor forever. In more details, our contribution was threefold: (i) We provide new complexity measures for communication efficiency of self-stabilizing protocols, especially in the stabilized phase or when there are no faults, (ii) On the negative side, we show that for non-trivial problems such as coloring, maximal matching, and maximal independent set, it is impossible to get (deterministic or probabilistic) self-stabilizing solutions where every participant communicates with less than every neighbor in the stabilized phase, and (iii) On the positive side, we present protocols for coloring, maximal matching, and maximal independent set such that a fraction of the participants communicates with exactly one neighbor in the stabilized phase.

The maximal matching problem has received considerable attention in the self-stabilizing community. Previous work has given different self-stabilizing algorithms that solves the problem for both the adversarial and fair distributed daemon, the sequential adversarial daemon, as well as the synchronous daemon. In [19] we presented a single self-stabilizing algorithm for this problem that unites all of these algorithms in that it stabilizes in the same number of moves as the previous best algorithms for the sequential adversarial, the distributed fair, and the synchronous daemon. In addition, the algorithm improves the previous best moves complexities for the distributed adversarial daemon from $O(n^2)$ and $O(\delta m)$ to $O(m)$ where n is the number of processes, m is the number of edges, and δ is the maximum degree in the graph.

In large scale multihop wireless networks, flat architectures are typically not scalable. Clustering was introduced to support self-organization and enable hierarchical routing. When dealing with multihop wireless networks, robustness is a crucial issue due to the dynamism of such networks. Several algorithms have been designed for clustering but to date, none of them has investigated the self-stabilization features of the resulting structure. In [21], we proved that a clustering algorithm that have previously exhibited good robustness properties, is actually self-stabilizing. We proposed several enhancements to the scheme to reduce the stabilization time and thus improve stability in a dynamic environment. The key technique to these enhancements is a localized self-stabilizing algorithm for Directed Acyclic Graph (DAG) construction. We provided extensive studies (both theoretical and experimental) that show that our approach enables efficient yet adaptive clustering in wireless multihop networks.

Unidirectional networks preclude many common techniques in self-stabilization from being used, such as preserving local predicates. In [28], we investigated the intrinsic complexity of achieving self-stabilization in unidirectional anonymous general networks, and focused on the classical vertex coloring problem. In more details, we proved a lower bound of n states per process (where n is the network size) and a recovery time of at least $n(n-1)/2$ actions in total. We also provided a deterministic algorithm with matching upper bounds that performs in arbitrary unidirectional anonymous graphs.

6.4. Extensions of Self-stabilization

Participant: Sébastien Tixeuil.

We tackled the open problem of snap-stabilization in message-passing systems [35]. Snap-stabilization is a nice approach to design protocols that withstand transient faults. Compared to the well-known self-stabilizing approach, snap-stabilization guarantees that the effect of faults is contained immediately after faults cease to occur. Our contribution is twofold: we show that (1) snap-stabilization is impossible for a wide class of problems if we consider networks with finite yet unbounded channel capacity; (2) snap-stabilization becomes possible in the same setting if we assume bounded-capacity channels. We propose three snap-stabilizing protocols working in fully-connected networks. Our work opens exciting new research perspectives, as it enables the snap-stabilizing paradigm to be implemented in actual networks.

The loop-free property provides interesting safety assurance in dynamic networks where edge-cost changes during operation of the protocol. The minimum spanning tree (MST) construction is a classical problem in Distributed Computing for creating a globally minimized structure distributedly. We presented [29] a new self-stabilizing MST protocol that improves on previous known approaches in several ways. First, it makes fewer system hypotheses as the size of the network (or an upper bound on the size) need not be known to the participants. Second, it is loop-free in the sense that it guarantees that a spanning tree structure is always preserved while edge costs change dynamically and the protocol adjusts to a new MST. Finally, time complexity matches the best known results, while space complexity results show that this protocol is the most efficient to date.

Distributed fault-tolerance can mask the effect of a limited number of permanent faults, while self-stabilization provides forward recovery after an arbitrary number of transient fault hit the system. FTSS protocols combine the best of both worlds since they are simultaneously fault-tolerant and self-stabilizing. To date, FTSS solutions either consider static (i.e. fixed point) tasks, or assume synchronous scheduling of the system components. We presented [39] the first study of dynamic tasks in asynchronous systems, considering the unison problem as a benchmark. Unison can be seen as a local clock synchronisation problem as neighbors must maintain digital clocks at most one time unit away from each other, and increment their own clock value infinitely often. We present many impossibility results for this difficult problem and propose a FTSS solution when the problem is solvable that exhibits optimal fault containment.

6.5. Byzantine Resilience in autonomous networks

Participant: Sébastien Tixeuil.

We studied [22] the problem of Byzantine-robust topology discovery in an arbitrary asynchronous network. We formally stated the weak and strong versions of the problem. The weak version requires that either each node discovers the topology of the network or at least one node detects the presence of a faulty node. The strong version requires that each node discovers the topology regardless of faults. We focused on non-cryptographic solutions to these problems. We explored their bounds. We proved that the weak topology discovery problem is solvable only if the connectivity of the network exceeds the number of faults in the system. Similarly, we showed that the strong version of the problem is solvable only if the network connectivity is more than twice the number of faults. We presented solutions to both versions of the problem. Our solutions match the established graph connectivity bounds. The programs are terminating, they do not require the individual nodes to know either the diameter or the size of the network. The message complexity of both programs is low polynomial with respect to the network size.

Given a set of robots with arbitrary initial location and no agreement on a global coordinate system, *convergence* requires that all robots asymptotically approach the exact same, but unknown beforehand, location. Robots are oblivious—they do not recall the past computations—and are allowed to move in a one-dimensional space. Additionally, robots cannot communicate directly, instead they obtain system related information only *via* visual sensors. We draw in [32] a connection between the convergence problem in robot networks, and the distributed *approximate agreement* problem (that requires correct processes to decide, for some constant ϵ , values distance ϵ apart and within the range of initial proposed values). Surprisingly, even though specifications are similar, the convergence implementation in robot networks requires specific assumptions about synchrony and Byzantine resilience. In more details, we proved necessary and sufficient conditions for the convergence of mobile robots despite a subset of them being Byzantine (*i.e.* they can exhibit arbitrary behavior). Additionally, we proposed a deterministic convergence algorithm for robot networks and analyze its correctness and complexity in various synchrony settings. The proposed algorithm tolerates f Byzantine robots for $(2f + 1)$ -sized robot networks in fully synchronous networks, $(3f + 1)$ -sized in semi-synchronous networks and $(4f + 1)$ -sized in asynchronous networks. The bounds obtained for the ATOM model are optimal for the class of *cautious* algorithms, which guarantee that correct robots always move inside the range of positions of the correct robots. We proposed in [33] the first deterministic algorithm that tolerates up to f Byzantine faults in $3f + 1$ -sized networks and performs in the asynchronous CORDA model. Our solution matches the previously established lower bound for the semi-synchronous ATOM model on the number of tolerated Byzantine robots. Our algorithm works under bounded scheduling assumptions for oblivious robots moving in a uni-dimensional space. We also studied [31] the convergence problem in fully asynchronous, uni-dimensional robot networks that are prone to Byzantine (*i.e.* malicious) failures. We proposed a deterministic algorithm that solves the problem in the most generic settings: fully asynchronous robots that operate in the non-atomic CORDA model. Our algorithm provides convergence in $5f + 1$ -sized networks where f is the upper bound on the number of Byzantine robots. Additionally, we proved that $5f + 1$ is a lower bound whenever robot scheduling is fully asynchronous. This contrasts with previous results in partially synchronous robots networks, where $3f + 1$ robots are necessary and sufficient.

6.6. Self-Stabilizing Control Infrastructure for HPC

Participants: Thomas Héroult, Camille Coti.

High performance computing platforms are becoming larger, leading to scalability and fault-tolerance issues for both applications and runtime environments (RTE) dedicated to run on such machines. After being deployed, usually following a spanning tree, a RTE needs to build its own communication infrastructure to manage and monitor the tasks of parallel applications. Previous works have demonstrated that the Binomial Graph topology (BMG) is a good candidate as a communication infrastructure for supporting scalable and fault-tolerant RTE.

In this work, we presented and analyzed a self-stabilizing algorithm to transform the underlying communication infrastructure provided by the launching service (usually a tree, due to its scalability during launch time) into a BMG, and maintain it in spite of failures. We demonstrated that this algorithm is scalable, tolerates transient failures, and adapts itself to topology changes.

The algorithms are scalable, in the sense that all process memory, number of established communication links, and size of messages are logarithmic with the number of elements in the system. The number of synchronous rounds to build the system is also logarithmic, and the number of asynchronous rounds in the worst case is square logarithmic with the number of elements in the system. Moreover, the self-stabilizing property of the algorithms presented induce fault-tolerance and self-adaptivity. Performance evaluation based on simulations predicts a fast convergence time (1/33s for 64K nodes), exhibiting the promising properties of such self-stabilizing approach.

We pursue this work by implementing and evaluating the algorithms in the STCI runtime environment to validate the theoretical results.

6.7. Large Scale Peer to Peer Performance Evaluations

Participant: Serge Petiton.

6.7.1. Large Scale Grid Computing

Recent progress has made possible to construct high performance distributed computing environments, such as computational grids and cluster of clusters, which provide access to large scale heterogeneous computational resources. Exploration of novel algorithms and evaluation of performance is a strategic research for the future of computational grid scientific computing for many important applications [111]. We adapted [88] an explicit restarted Lanczos algorithm on a world-wide heterogeneous grid platform. This method computes one or few eigenpairs of a large sparse real symmetric matrix. We take the specificities of computational resources into account and deal with communications over the Internet by means of techniques such as out-of-core and data persistence. We also show that a restarted algorithm and the combination of several paradigms of parallelism are interesting in this context. We perform many experimentations using several parameters related to the Lanczos method and the configuration of the platform. Depending on the number of computed Ritz eigenpairs, the results underline how critical the choice of the dimension of the working subspace is. Moreover, the size of platform has to be scaled to the order of the eigenproblem because of communications over the Internet.

6.7.2. High Performance Cluster Computing

Grid computing focuses on making use of a very large amount of resources from a large-scale computing environment. It intends to deliver high-performance computing over distributed platforms for computation and data-intensive applications. We propose [130] an effective parallel hybrid asynchronous method to solve large sparse linear systems by the use of a Grid Computing platform Grid5000. This hybrid method combines a parallel GMRES(m) (Generalized Minimum RESidual) algorithm with the Least Square method that needs some eigenvalues obtained from a parallel Arnoldi algorithm. All of these algorithms run on the different processors of the platform Grid5000. Grid5000, a 5000 CPUs nation-wide infrastructure for research in Grid computing, is designed to provide a scientific tool for computing. We discuss the performances of this hybrid method deployed on Grid5000, and compare these performances with those on the IBM SP series supercomputers.

6.7.3. Large Scale Power aware Computing

Energy conservation is a dynamic topic of research in High Performance Computing and Cluster Computing. Power-aware computing for heterogeneous world-wide Grid is a new track of research. We have studied and evaluated the impact of the heterogeneity of the computing nodes of a Grid platform on the energy consumption. We propose to take advantage of the slack-time caused by the heterogeneity in order to save energy with no significant loss of performance by using Dynamic Voltage Scaling (DVS) in a distributed eigensolver [89]. We show that using DVS only during the slack-time does not penalize the performances but it does not provide significant energy savings. If DVS is applied to all the execution, we get important global and local energy savings (respectively up to 9% and 20%) without a significant rise of the wall-clock times.

6.8. High Performance Linear Algebra on the Grid

Participants: Thomas Héroult, Camille Coti.

Previous studies have reported that common dense linear algebra operations do not achieve speed up by using multiple geographical sites of a computational grid. Because such operations are the building blocks of most scientific applications, conventional supercomputers are still strongly predominant in high-performance computing and the use of grids for speeding up large-scale scientific problems is limited to applications exhibiting parallelism at a higher level.

In this work, we have identified two performance bottlenecks in the distributed memory algorithms implemented in ScaLAPACK, a state-of-the-art dense linear algebra library. First, because ScaLAPACK assumes a homogeneous communication network, the implementations of ScaLAPACK algorithms lack locality in their communication pattern. Second, the number of messages sent in the ScaLAPACK algorithms is significantly greater than other algorithms that trade flops for communication.

This year, we presented a new approach for computing a QR factorization one of the main dense linear algebra kernels of tall and skinny matrices in a grid computing environment that overcomes these two bottlenecks. Our contribution is to articulate a recently proposed algorithm (Communication Avoiding QR) with a topology-aware middleware (QCG-OMPI) in order to confine intensive communications (ScaLAPACK calls) within the different geographical sites.

An experimental study conducted on the Grid5000 platform shows that the resulting performance increases linearly with the number of geographical sites on large-scale problems (and is in particular consistently higher than ScaLAPACKs).

6.9. Emulation of Volatile Systems

Participants: Thomas Largillier, Benjamin Quetier, Sylvain Peyronnet, Thomas Hérault, Franck Cappello.

In the process of developing grid applications, people need to often evaluate the robustness of their work. Two common approaches are simulation, where one can evaluate his software and predict behaviors under conditions usually unachievable in a laboratory experiment, and experimentation, where the actual application is launched on an actual grid. However simulation could ignore unpredictable behaviors due to the abstraction done and experimentation does not guarantee a controlled and reproducible environment, and simulation often introduces a high level of abstraction that make the discovery and study of unexpected, but real, behaviors a rare event.

In this work, we proposed an emulation platform for parallel and distributed systems including grids where both the machines and the network are virtualized at a low level. The use of virtual machines allows us to test highly accurate failure injection since we can destroy virtual machines, and network virtualization provides low-level network emulation. Failure accuracy is a criteria that evaluates how realistic a fault is. The accuracy of our framework has been evaluated through a set of micro benchmarks and a very stable P2P system called Pastry.

We are in the process of developing a fault injection tool to work with the platform. it will be an extension of the work started in the tool Fail. The interest of this work is that using Xen virtual machines will allow to model strong adversaries since it is possible to have virtual machines with shared memory. These adversaries will be stronger since they will be able to use global fault injection strategies.

6.10. Exascale Systems

Participant: Franck Cappello.

Over the last 20 years, the open-source community has provided more and more software on which the world's high-performance computing systems depend for performance and productivity. The community has invested millions of dollars and years of effort to build key components. Although the investments in these separate software elements have been tremendously valuable, a great deal of productivity has also been lost because of the lack of planning, coordination, and key integration of technologies necessary to make them work together smoothly and efficiently, both within individual petascale systems and between different systems. A repository gatekeeper and an email discussion list can coordinate open-source development within a single project, but there is no global mechanism working across the community to identify critical holes in the overall software environment, spot opportunities for beneficial integration, or specify requirements for more careful coordination. It seems clear that this completely uncoordinated development model will not provide the software needed to support the unprecedented parallelism required for peta/exascale computation on millions of cores, or the flexibility required to exploit new hardware models and features, such as transactional memory, speculative execution, and GPUs. We presented in [17] a rational promoting that the community must work together to prepare for the challenges of exascale computing, ultimately combing their efforts in a coordinated International Exascale Software Project.

Over the past few years resilience has become a major issue for high-performance computing (HPC) systems, in particular in the perspective of large petascale systems and future exascale systems. These systems will typically gather from half a million to several millions of central processing unit (CPU) cores running up to a billion threads. From the current knowledge and observations of existing large systems, it is anticipated that exascale systems will experience various kind of faults many times per day. It is also anticipated that the current approach for resilience, which relies on automatic or application level checkpoint/restart, will not work because the time for checkpointing and restarting will exceed the mean time to failure of a full system. This set of projections leaves the community of fault tolerance for HPC systems with a difficult challenge: finding new approaches, which are possibly radically disruptive, to run applications until their normal termination, despite the essentially unstable nature of exascale systems. Yet, the community has only five to six years to solve the problem. In order to start addressing this challenge, we synthesized in [14] the motivations, observations and research issues considered as determinant of several complimentary experts of HPC in applications, programming models, distributed systems and system management.

As a first step to address the resilience challenge, we conducted a comprehensive study of the state of the art published in [13]. The emergence of petascale systems and the promise of future exascale systems have reinvigorated the community interest in how to manage failures in such systems and ensure that large applications, lasting several hours or tens of hours, are completed successfully. Most of the existing results for several key mechanisms associated with fault tolerance in high-performance computing (HPC) platforms follow the rollback-recovery approach. Over the last decade, these mechanisms have received a lot of attention from the community with different levels of success. Unfortunately, despite their high degree of optimization, existing approaches do not fit well with the challenging evolutions of large-scale systems. There is room and even a need for new approaches. Opportunities may come from different origins: diskless checkpointing, algorithmic-based fault tolerance, proactive operation, speculative execution, software transactional memory, forward recovery, etc. We provided the following contributions: (1) we summarize and analyze the existing results concerning the failures in large-scale computers and point out the urgent need for drastic improvements or disruptive approaches for fault tolerance in these systems; (2) we sketch most of the known opportunities and analyze their associated limitations; (3) we extract and express the challenges that the HPC community will have to face for addressing the stringent issue of failures in HPC systems.

6.11. High performance scientific computing

Participants: Laura Grigori, Guy Atenekeng, Simplice Donfack, Amal Khabou, Pawan Kumar, Federico Stivoli, Ke Wang.

The focus of this research is on the design of efficient parallel algorithms for solving problems in numerical linear algebra, as solving very large sets of linear equations and large least squares problems, often with millions of rows and columns. These problems arise in many numerical simulations, and solving them is very time consuming.

6.11.1. Communication avoiding algorithms for LU and QR factorizations

This research focuses on developing new algorithms for linear algebra problems, that minimize the required communication, in terms of both latency and bandwidth. The results we have obtained to date concern algorithms for solving dense linear systems and dense least squares problems.

In [61] we study algorithms for performing the LU and QR factorizations of dense matrices. Recently, two communication optimal algorithms have been introduced for distributed memory architectures, referred to as communication avoiding CALU and CAQR (joint work with J. Demmel and M. Hoemmen from U.C. Berkeley, J. Langou from C.U. Denver, and H. Xiang from University Paris 6). We study two algorithms based on CAQR and CALU that are adapted to multicore architectures. They combine ideas to reduce communication from communication avoiding algorithms with asynchronism and dynamic task scheduling. For matrices that are tall and skinny, that is, they have many more rows than columns, the two algorithms outperform the corresponding algorithms from Intel MKL vendor library on a dual-socket, quad-core machine based on Intel Xeon EMT64 processor and on a four-socket, quad-core machine based on AMD Opteron processor. For

matrices of size $m = 10^5$ and $m = 10^6$, for n varying from 10 to 1000, multithreaded CALU outperforms the corresponding routine dgetrf from Intel MKL library up to a factor of 2.3 and the corresponding routine dgetrf from AMD (ACML library) up to a factor of 5. Multithreaded CAQR outperforms by a factor of 5.3 the corresponding dgeqrf routine from MKL library. For square matrices, CALU is slightly slower than MKL dgetrf routine when $m = n < 5000$, while for $m = n = 10^4$ it is slightly faster than this routine up to a factor of 1.5 and CAQR is less performant than the corresponding routine in the other libraries.

In 2008 we have introduced a Communication-Avoiding LU factorization (CALU) algorithm for computing in parallel the LU factorization of a dense matrix A distributed in a two-dimensional (2D) layout. To decrease the communication required in the LU factorization, CALU uses a new pivoting strategy, referred to as ca-pivoting, that may lead to a different row permutation than the classic LU factorization with partial pivoting. We have further investigated the numerical stability of CALU. Our numerical results show that ca-pivoting scheme is stable in practice. We observe that it behaves as a threshold pivoting, and in practical experiments $|L|$ is bounded by 3, while in LU factorization with partial pivoting, $|L|$ is bounded by 1, where $|L|$ denotes the matrix of absolute values of the entries of L . Extensive testing on many different matrices always resulted in residuals $\|Ax - b\|$ comparable to those from conventional partial pivoting. In particular we have shown that CALU is equivalent to performing GEPP on a larger matrix formed by entries of the input matrix (sometimes slightly perturbed) and zeros. The paper describing this research is in preparation.

6.11.2. Preconditioning techniques

A different direction of research is related to preconditioning large sparse linear systems of equations. In this research we consider different preconditioners based on incomplete factorizations. The tangential filtering is an incomplete factorization technique where it is possible to ensure that the factorization will coincide with the original matrix for some specified vector. This research is performed in the context of ANR PETAL project. The participants at INRIA Saclay are L. Grigori, P. Kumar and K. Wang.

Recent research has shown that ILU combined with tangential filtering leads to very efficient preconditioner for matrices arising from the discretization of scalar equations on structured grids and in the previous year we have investigated further their properties.

The problem of solving block tridiagonal linear systems arising from the discretization of PDE is considered. The nested factorization preconditioner introduced by [J. R. Appleyard and I. M. Cheshire, *Nested Factorization*, SPE 12264, presented at the Seventh SPE Symposium on Reservoir Simulation, San Francisco, 1983] is an effective preconditioner for certain class of problems and a similar method is implemented in Schlumberger's Eclipse oil reservoir simulator. In [63], a relaxed version of Nested Factorization preconditioner is proposed as a replacement to ILU0. Indeed the proposed preconditioner is SPD and leads to a stable splitting if the input matrix is S.P.D. . For ILU0, equivalent properties hold if the input matrix is a M-matrix. Moreover it has no storage cost. Effective multiplicative/additive preconditioning is achieved in combination with Tangential filtering preconditioner with the filter vector chosen as vector of *ones*. Numerical tests are carried out with both additive and multiplicative combinations. With this setup the new preconditioner is as robust as the combination of ILU0 with tangential filtering preconditioner.

We also designed preconditioners based on Kronecker product approximation or Schilder factorization for saddle point problems arising from PDE or optimization [23].

7. Other Grants and Activities

7.1. Regional, National and International Actions

7.1.1. Activities starting in 2009

- Franck Cappello, Co-Director of the **INRIA - Illinois Joint Laboratory** on PetaScale Computing, since 2009

7.1.2. Other activities

- **ANR SPADES** Coordinated by LIP-ENS Lyon. (Sylvain Peyronnet, Franck Cappelto)
- **Défi ANR SECSI** Participant to this challenge. From September 2008 to August 2010. Managed by the SAIC. (Thomas Hérault, Sylvain Peyronnet, Sébastien Tixeuil)
- **ANR Cosinus project MIDAS - Microwave Data Analysis for petaScale computers** December 2009 - December 2012 (http://www.apc.univ-paris7.fr/APC_CS/Recherche/Adamis/MIDAS09/). Collaboration with APC, University Paris 7 and Lawrence Berkeley Laboratory. This is an interdisciplinary project devised to bring together cosmologists, computational physicists, computer scientists and applied mathematicians to face the challenge of the tremendous volume of data as anticipated from current and forthcoming Cosmic Microwave Background (CMB) experiments. (Laura Grigori, Coordinator for INRIA Saclay, F. Cappelto, J. Falcou, T. Hérault, S. Peyronnet)
- **ANR Cosinus project PETAL- PETascale ALgorithms for preconditioning for scientific applications** January 2009- December 2010. Collaboration with Laboratoire Lions - Université 6, IFP, INRIA Bordeaux and CEA, UC Berkeley and Argonne. The goal is to investigate preconditioning techniques on multicore architectures and apply them on real world applications from IFP, CEA and Argonne. (Laura Grigori, Principal Investigator)
- **Digitéo DIM-08 project X-Scale-NL – Scheduling and numerical libraries enabling scientific applications on petascale machines** 2008-2011. Funding for a Phd student and travel (114000 euros). Participants: Laura Grigori (Principal Investigator), F. Cappelto (INRIA), T. Hérault, S. Peyronnet (LRI) and two foreign collaborators: J. Demmel from UC Berkeley and J. Darbon from UC Los Angeles.
- **PECO-NEI RFR** with Eastern Europe Countries, 2006 - 2009, PI L. Grigori
- INRIA Associated Team "**F-J Grid**" with University of Tsukuba, head: Franck Cappelto
- INRIA funding, **MPI-V**, collaboration with UTK, LALN and ANL, head: Franck Cappelto
- Regional Council "**Grid eXplorer**", 3 years (2006-2009), co-chair: Franck Cappelto
- ANR Jeunes chercheurs **DSL-Lab**: G. Fedak, 3 years (2005-2008)
- ANR CIS Project **FF2A3**, 3 years (2007 - 2010), PI F. Hecht, subproject head L. Grigori
- European project. **Grid4All**, 3 years (2006-2009), management board representative: Franck Cappelto, task head: Gilles Fedak
- **CARRIOCAS**, Pole de Compétitivité System@tic, 3 years (2006-2009), <http://www.carriocas.org/>, Franck Cappelto
- **HipCal**, ANR CIS, 3 years (2006-2009), <http://hipcal.lri.fr/wakka.php?wiki=PagePrincipale>, Franck Cappelto
- **QosCosGrid**, European Project, 09/2006 to 05/2009, Thomas Hérault
- participation to **OpenMPI consortium** since 01/2008, Thomas Hérault

8. Dissemination

8.1. Services to the Scientific Community

8.1.1. Research Administration

- Brigitte Rozoy, past position of University Vice-President and position at the research ministry.

8.1.2. Book/Journal edition

- Joffroy Beauquier, PPL, Parallel Processing Letters

- Franck Cappello, Cluster Computing Journal, Springer, Netherlands, since 2006
- Franck Cappello, Journal of Grid Computing, Springer Netherlands, since 2003
- Franck Cappello, Journal of Grid and utility computing, Inderscience, since 2004

8.1.3. Conference Organisation

- Franck Cappello, Program co-chair, IEEE CCGRID'2009
- Franck Cappello, Area co-chair, IEEE-ACM Supercomputing '09
- Franck Cappello, General co-chair, IPDPS/PCGrid'09

8.1.4. Editorial Committee membership

- Sébastien Tixeuil, Technique et Science Informatiques, since 2005

8.1.5. Steering Committee membership

- Franck Cappello, IEEE/ACM HPDC
- Franck Cappello, IEEE/ACM CCGRID
- Brigitte Rozoy, MSR, Modélisation des Systèmes Réactifs, 2006-2009

8.1.6. Program Committee membership

- Franck Cappello, IEEE HPDC'2009
- Franck Cappello, IEEE/ACM SC'2009
- Franck Cappello, IEEE/ACM CCGRID'2009
- Franck Cappello, IEEE/ACM Grid 2009
- Franck Cappello, SIMUTOOLS'09
- Franck Cappello, ICPP 2009
- Franck Cappello, SSS 2009
- Franck Cappello, first workshop on Automated Control for Data Centers and Clouds (ACDC) 2009
- Franck Cappello, 3rd Workshop on Virtualization Technology in Distributed Computing (VTDC) 2009
- Franck Cappello, Resilience 2009
- Franck Cappello, ISPA'09 (the 7th IEEE International Symposium on Parallel and Distributed Processing with Applications)
- Franck Cappello, EuroPVM/MPI 2009
- Franck Cappello, Workshop on Resiliency for Petascale HPC 2009
- Laura Grigori, Supercomputing 2009
- Laura Grigori, International Conference on Parallel Processing and Applied Mathematics PPAM09
- Sylvain Peyronnet, ISVC 2009
- Sébastien Tixeuil, WRAS 2009, Program co-chair
- Sébastien Tixeuil, INTERNET 2009
- Sébastien Tixeuil, OPODIS 2009
- Sébastien Tixeuil, SSS 2009
- Sébastien Tixeuil, ISPS 2009
- Sébastien Tixeuil, CODS 2009

9. Bibliography

Major publications by the team in recent years

- [1] R. BOLZE, F. CAPPELLO, E. CARON, M. J. DAYDÉ, F. DESPREZ, E. JEANNOT, Y. JÉGOU, S. LANTERI, J. LEDUC, N. MELAB, G. MORNET, R. NAMYST, P. PRIMET, B. QUÉTIER, O. RICHARD, E.-G. TALBI, T. IRENA. *Grid'5000: a large scale and highly reconfigurable experimental Grid testbed.*, in "International Journal of High Performance Computing Applications", vol. 20, n^o 4, November 2006, p. 481-494.
- [2] A. BOUTEILLER, T. HÉRAULT, G. KRAWEZIK, P. LEMARINIER, F. CAPPELLO. *MPICH-V Project: a Multiprotocol Automatic Fault Tolerant MPI*, in "International Journal of High Performance Computing Applications", vol. 20, n^o 3, 2005, p. 319-333.
- [3] F. CAPPELLO, S. DJILALI, G. FEDAK, T. HÉRAULT, F. MAGNIETTE, V. NÉRI, O. LODYGENSKY. *Computing on Large Scale Distributed Systems: XtremWeb Architecture, Programming Models, Security, Tests and Convergence with Grid*, in "FGCS Future Generation Computer Science", 2004.
- [4] T. HERMAN, S. TIXEUIL. *A Distributed TDMA Slot Assignment Algorithm for Wireless Sensor Networks*, in "Proceedings of the First Workshop on Algorithmic Aspects of Wireless Sensor Networks (AlgoSensors'2004), Turku, Finland", S. NIKOLETSEAS, J. D. ROLIM (editors), Lecture Notes in Computer Science, vol. 3121, Springer-Verlag, July 2004, p. 45-58, <http://dx.doi.org/10.1007/b98740>.
- [5] W. HOARAU, S. TIXEUIL, F. VAUCHELLES. *FAIL-FCI: Versatile Fault-Injection*, in "Future Generation Computer Systems", vol. 23, n^o 7, 2007, p. 913-919.
- [6] T. HÉRAULT, R. LASSAIGNE, S. PEYRONNET. *APMC 3.0: Approximate Verification of Discrete and Continuous Time Markov Chains*, in "Proceedings of the 3rd International Conference on the Quantitative Evaluation of SysTems (QEST'06), California, USA", September 2006.
- [7] C. JOHNEN, L. O. ALIMA, A. K. DATTA, S. TIXEUIL. *Optimal Snap-stabilizing Neighborhood Synchronizer in Tree Networks*, in "Parallel Processing Letters", vol. 12, n^o 3 & 4, 2002, p. 327-340.
- [8] T. MASUZAWA, S. TIXEUIL. *On Bootstrapping Topology Knowledge in Anonymous Networks*, in "Eighth International Symposium on Stabilization, Safety, and Security on Distributed Systems (SSS 2006), Dallas, Texas", A. K. DATTA, M. GRADINARIU (editors), Lecture Notes in Computer Science, vol. 4, Springer Verlag, November 2006, p. 454-468, <http://doi.acm.org/10.1145/1462187.1462195>.
- [9] B. WEI, G. FEDAK, F. CAPPELLO. *Scheduling Independent Tasks Sharing Large Data Distributed with BitTorrent*, in "IEEE/ACM Grid'2005 workshop Seattle, USA", 2005.

Year Publications

Doctoral Dissertations and Habilitation Theses

- [10] J. CLÉMENT. *Algorithmique probabiliste pour systèmes distribués émergents*, Université Paris-Sud, Orsay, France, October 19th 2009, Ph. D. Thesis.

- [11] C. COTI. *Environnements d'exécution pour applications parallèles communiquant par passage de messages pour les systèmes à grande échelle et les grilles calcul*, Université Paris Sud-XI, November 2009, Ph. D. Thesis.
- [12] A. REZMERITA. *Contribution aux intergiciels et protocoles pour les grappes virtuelles*, Université Paris-Sud, Orsay, France, September 28 2009, Ph. D. Thesis.

Articles in International Peer-Reviewed Journal

- [13] F. CAPPELLO. *Fault Tolerance in Petascale/ Exascale Systems: Current Knowledge, Challenges and Research Opportunities*, in "International Journal of High Performance Computing Applications", vol. 23, n^o 3, 2009, p. 212–226, <http://dx.doi.org/10.1177/1094342009106189>.
- [14] F. CAPPELLO, A. GEIST, B. GROPP, L. KALE, B. KRAMER, M. SNIR. *Toward Exascale Resilience*, in "International Journal of High Performance Computing Applications", vol. 23, n^o 4, November 2009, p. 374–388, <http://dx.doi.org/10.1177/1094342009347767>us.
- [15] L. CHOY, S. G. PETITON, M. SATO. *Resolution of large symmetric eigenproblems on a world-wide grid*, in "International Journal of Grid and Utility Computing", vol. 1, n^o 2, 2009, p. 71–85, <http://inderscience.metapress.com/openurl.asp?genre=article&eissn=1741-8488&volume=1&issue=2&spage=71>.
- [16] P. DANTURI, M. NESTERENKO, S. TIXEUIL. *Self-stabilizing Philosophers with Generic Conflicts*, in "ACM Transactions of Adaptive and Autonomous Systems (TAAS)", vol. 4, January 2009, p. 1–20, <http://doi.acm.org/10.1145/1462187.1462194>.
- [17] J. DONGARRA, P. BECKMAN, P. AERTS, F. CAPPELLO, T. LIPPERT, S. MATSUOKA, P. MESSINA, T. MOORE, R. STEVENS, A. TREFETHEN, M. VALERO. *The International Exascale Software Project: a Call To Cooperative Action By the Global High-Performance Community*, in "International Journal of High Performance Computing Applications", vol. 23, n^o 4, 2009, p. 309–322, <http://dx.doi.org/10.1177/1094342009347714>usnldejpuks.
- [18] G. FEDAK, H. HE, F. CAPPELLO. *BitDew: A data management and distribution service with multi-protocol file transfer and metadata abstraction*, in "Journal of Network and Computer Applications", vol. 32, n^o 5, 2009, p. 961 - 975, <http://www.sciencedirect.com/science/article/B6WKB-4W38RND-1/2/3ba4f5eab1c505526f1b2084c1e967dd>, Next Generation Content Networks.
- [19] F. MANNE, M. MJELDE, L. PILARD, S. TIXEUIL. *A New Self-Stabilizing Maximal Matching Algorithm*, in "Theoretical Computer Science (TCS)", vol. 410, 2009, p. 1336-1345, <http://dx.doi.org/10.1016/j.tcs.2008.12.022>.
- [20] T. MASUZAWA, S. TIXEUIL. *On Bootstrapping Topology Knowledge in Anonymous Networks*, in "ACM Transactions on Adaptive and Autonomous Systems (TAAS)", vol. 4, n^o 1, 2009, p. 1–27, <http://doi.acm.org/10.1145/1462187.1462195>.
- [21] N. MITTON, B. SÉRICOLA, S. TIXEUIL, E. FLEURY, I. GUÉRIN-LASSOUS. *Self-stabilization in Self-organized Multihop Wireless Networks*, in "Ad hoc and Sensor Wireless Networks", 2009.

- [22] M. NESTERENKO, S. TIXEUIL. *Discovering Network Topology in the Presence of Byzantine Nodes*, in "IEEE Transactions on Parallel and Distributed Systems", vol. 20, n^o 12, October 2009, p. 1777-1789, <http://doi.ieeecomputersociety.org/10.1109/TPDS.2009.25>.
- [23] H. XIANG, L. GRIGORI. *Kronecker product approximation preconditioners for convection-diffusion model problems*, in "Numerical Linear Algebra with Applications", 2009, <http://dx.doi.org/10.1002/nla.666>.
- International Peer-Reviewed Conference/Proceedings**
- [24] E. AGULLO, C. COTI, J. DONGARRA, T. HÉRAULT, J. LANGOU. *QR Factorization of Tall and Skinny Matrices in a Grid Computing Environment*, in "24th IEEE International Parallel & Distributed Processing Symposium (IPDPS'10), Atlanta, Ga", April 2010, to appear.
- [25] P. BAR, C. COTI, D. GROEN, T. HÉRAULT, V. KRAVTSOV, A. SCHUSTER, M. SWAIN. *Running parallel applications with topology-aware grid middleware*, in "5th IEEE International Conference on e-Science (eScience 2009)", December 2009, to appear US .
- [26] J. BEAUQUIER, J. BURMAN, J. CLÉMENT, S. KUTTEN. *Brief announcement: non-self-stabilizing and self-stabilizing gathering in networks of mobile agents – the notion of speed*, in "PODC '09: Proceedings of the 28th ACM symposium on Principles of distributed computing", ACM, August 10-12 2009, p. 286-287, <http://doi.acm.org/10.1145/1582716.1582768il>.
- [27] J. BEAUQUIER, J. BURMAN, S. KUTTEN. *Making Population Protocols Self-stabilizing*, in "Proceedings of 11th International Symposium on Stabilization, Safety, and Security of Distributed Systems, SSS 2009, Lyon, France", vol. 5873, November 3-6 2009, p. 90-104, http://dx.doi.org/10.1007/978-3-642-05118-0_7il.
- [28] S. BERNARD, S. DEVISMES, M. G. POTOP-BUTUCARU, S. TIXEUIL. *Optimal Deterministic Self-stabilizing Vertex Coloring in Unidirectional Anonymous Networks*, in "Proceedings of the IEEE International Conference on Parallel and Distributed Processing Systems (IPDPS 2009), Rome, Italy", IEEE Press, May 23-29 2009, <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2009.5161053>.
- [29] L. BLIN, M. G. POTOP-BUTUCARU, S. ROVEDAKIS, S. TIXEUIL. *A New Self-Stabilizing Minimum Spanning Tree Construction with Loop-free Property*, in "Proceedings of DISC 2009, Elche, Spain", Lecture Notes in Computer Science, Springer Berlin / Heidelberg, September 2009, <http://arxiv.org/abs/0905.2287>.
- [30] G. BOSILCA, C. COTI, T. HÉRAULT, P. LEMARINIER, J. DONGARRA. *Constructing Resilient Communication Infrastructure for Runtime Environments*, in "International Conference in Parallel Computing (ParCo2009), Lyon, France", September 2009.
- [31] Z. BOUZID, M. G. POTOP-BUTUCARU, S. TIXEUIL. *Byzantine Convergence in Robots Networks: The Price of Asynchrony*, in "Proceedings of OPODIS 2009, Nîmes, France", Lecture Notes in Computer Science, Springer Berlin / Heidelberg, December 2009, <http://arxiv.org/abs/0908.0390>.
- [32] Z. BOUZID, M. G. POTOP-BUTUCARU, S. TIXEUIL. *Byzantine-resilient Convergence in Oblivious Robot Networks*, in "International Conference on Distributed Systems and Networks (ICDCN 2009), Hyderabad, India", vol. 5408, Springer Berlin / Heidelberg, January 3-6 2009, p. 275-280, http://dx.doi.org/10.1007/978-3-540-92295-7_33.

- [33] Z. BOUZID, M. G. POTOP-BUTUCARU, S. TIXEUIL. *Optimal Byzantine Resilient Convergence in Asynchronous Robot Networks*, in "Proceedings of SSS 2009., Lyon, France", Lecture Notes in Computer Science, Springer Berlin / Heidelberg, November 2009, <http://arxiv.org/abs/0906.0651>.
- [34] C. COTI, T. HÉRAULT, F. CAPPELLO. *MPI Applications on Grids: a Topology-Aware Approach*, in "Proceedings of the 15th European Conference on Parallel and Distributed Computing (EuroPar'09), Delft, the Netherlands", Lecture Notes in Computer Science, vol. 5704, Springer Berlin / Heidelberg, August 2009, p. 466-477, http://dx.doi.org/10.1007/978-3-642-03869-3_45.
- [35] S. DELAËT, S. DEVISMES, M. NESTERENKO, S. TIXEUIL. *Snap-Stabilization in Message-Passing Systems*, in "International Conference on Distributed Systems and Networks (ICDCN 2009)", LNCS, n^o 5404, January 2009, p. 281-286, http://dx.doi.org/10.1007/978-3-540-92295-7_34.
- [36] S. DEVISMES, T. MASUZAWA, S. TIXEUIL. *Communication Efficiency in Self-Stabilizing Silent Protocols*, in "Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS 2009), Montreal, Canada", IEEE Press, June 2009, <http://hal.inria.fr/inria-00340805/fr/>.
- [37] S. DEVISMES, F. PETIT, S. TIXEUIL. *Optimal Probabilistic Ring Exploration by Asynchronous Oblivious Robots*, in "Proceedings of 16th International Colloquium, SIROCCO 2009, Piran, Slovenia", Lecture Notes in Computer Science, vol. 5869, Springer-Verlag Berlin Heidelberg, May 25-27 2009, <http://hal.inria.fr/inria-00360305/fr/>.
- [38] S. DONFACK, L. GRIGORI, A. K. GUPTA. *Implementing communication-avoiding LU and QR factorizations for tall and skinny matrices on multicore architectures*, in "Proceedings of 24th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2010), Atlanta, USA", April 19-23 2010.
- [39] S. DUBOIS, M. G. POTOP-BUTUCARU, S. TIXEUIL. *Brief Announcement: Dynamic FTSS in Asynchronous Systems: the Case of Unison*, in "Proceedings of DISC 2009, Elche, Spain", Lecture Notes in Computer Science, Springer Berlin / Heidelberg, September 2009, <http://arxiv.org/abs/0904.4615>.
- [40] H. HE, G. FEDAK, B. TANG, F. CAPPELLO. *BLAST Application with Data-Aware Desktop Grid Middleware*, in "CCGRID'09: Proceedings of the 2009 9th IEEE/ACM International Symposium on Cluster Computing and the Grid", IEEE Computer Society, May 2009, p. 284-291, <http://dx.doi.org/10.1109/CCGRID.2009.91>.
- [41] T. HÉRAULT, T. LARGILLIER, S. PEYRONNET, B. QUÉTIÉ, F. CAPPELLO, M. JAN. *High accuracy failure injection in parallel and distributed systems using virtualization*, in "Proceedings of the 6th Conference on Computing Frontiers, Ischia, Italy", 2009, p. 193-196, <http://doi.acm.org/10.1145/1531743.1531774>.
- [42] D. KONDO, B. JAVADI, P. MALÉCOT, F. CAPPELLO, D. P. ANDERSON. *Cost-benefit analysis of Cloud Computing versus desktop grids*, in "IEEE International Symposium on Parallel & Distributed Processing 2009, Rome, Italy", IEEE Computer Society, May 23-29 2009, p. 1-12, <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2009.5160911>.
- [43] R. PODESTA, V. INIESTA, A. REZMERITA, F. CAPPELLO. *An Information Brokering Service Provider (IBSP) for Virtual Clusters*, in "On the Move to Meaningful Internet Systems: OTM 2009, Lecture Notes in Computer Science, Vilamoura, Portugal", vol. 5870/2009, Springer Berlin / Heidelberg, November 1-6 2009, p. 165-182, http://dx.doi.org/10.1007/978-3-642-05148-7_13.

- [44] L. SHANG, S. G. PETITON, N. EMAD, X. YANG, Z. WANG. *Extending YML to Be a Middleware for Scientific Cloud Computing*, in "Cloud Computing, First International Conference, CloudCom 2009, Beijing, China", December 1-4 2009, p. 662-667, http://dx.doi.org/10.1007/978-3-642-10665-1_69.
- [45] L. SHANG, S. G. PETITON, M. HUGUES. *A New Parallel Paradigm for Block-Based Gauss-Jordan Algorithm*, in "Eighth International Conference on Grid and Cooperative Computing 2009, Los Alamitos, CA, USA", IEEE Computer Society, August 27-29 2009, p. 193-200, <http://doi.ieeecomputersociety.org/10.1109/GCC.2009.75>.
- [46] P. SOLLICH, M. URRY, C. COTI. *Kernels and learning curves for Gaussian process regression on random graphs*, in "Advances in Neural Information Processing Systems 23 (NIPS 2009), Vancouver, BC, Canada", December 2009, p. 1723-1731, http://books.nips.cc/papers/files/nips22/NIPS2009_0092.pdf.

National Peer-Reviewed Conference/Proceedings

- [47] S. BERNARD, S. DEVISMES, K. PAROUX, M. G. POTOP-BUTUCARU, S. TIXEUIL. *Sur le Coloriage Auto-stabilisant dans les Réseaux Unidirectionnels Anonymes*, in "Proceedings of Algotel 2009", May 2009, <http://hal.inria.fr/inria-00384649/en/>.
- [48] S. DELAËT, S. DEVISMES, M. NESTERENKO, S. TIXEUIL. *Stabilisation Instantanée dans les Systèmes à Passage de Messages*, in "Proceedings of Algotel 2009", May 2009, <http://hal.archives-ouvertes.fr/inria-00383350/en/>.
- [49] S. DEVISMES, F. PETIT, S. TIXEUIL. *Exploration Optimale Probabiliste d'un Anneau par des Robots Semi-Synchrones et Amnésiques*, in "Proceedings of Algotel 2009, Carry-Le-Rouet, France", May 2009, <http://hal.archives-ouvertes.fr/inria-00383351/en/>.

Scientific Books (or Scientific Book chapters)

- [50] F. CAPPELLO, G. FEDAK, D. KONDO, P. MALÉCOT, A. REZMERITA. *Desktop Grids: From Volunteer Distributed Computing to High Throughput Computing Production Platforms*, in "Handbook of Research on Scalable Computing Technologies", IGI Global, 2009, p. 31-61, <http://dx.doi.org/10.4018/978-1-60566-661-7.ch003>.
- [51] T. HÉRAULT, M. JAN, T. LARGILLIER, S. PEYRONNET, B. QUÉTIER, F. CAPPELLO. *Emulation platform for high accuracy failure injection in grids*, in "High Performance & Large Scale Computing", vol. 18, IOS Press, 2009, <http://dx.doi.org/0.3233/978-1-60750-073-5-127>.
- [52] S. TIXEUIL. *Self-stabilizing Algorithms*, in "Algorithms and Theory of Computation Handbook, Second Edition", Chapman & Hall/CRC Applied Algorithms and Data Structures, Taylor & Francis, 2009.

Books or Proceedings Editing

- [53] F. CAPPELLO, T. HÉRAULT, J. DONGARRA (editors). *Selected papers from the 14th European PVM/MPI Users Group Meeting, Special Issue of Parallel Computing Journal*, vol. 35, n° 12, Elsevier, December 2009, <http://dx.doi.org/10.1016/j.parco.2009.11.001>.
- [54] H. JIN, C.-H. HSU, F. CAPPELLO (editors). *Special Issue on "Peer-to-Peer Grid Technologies", Future Generation Computer Systems journal*, Elsevier, 2009.

Research Reports

- [55] L. BLIN, M. G. POTOP-BUTUCARU, S. ROVEDAKIS, S. TIXEUIL. *A New Self-Stabilizing Minimum Spanning Tree Construction with Loop-free Property*, arXiv, May 2009, <http://arxiv.org/abs/0905.2287>, Technical report.
- [56] Z. BOUZID, M. G. POTOP-BUTUCARU, S. TIXEUIL. *Byzantine Convergence in Robots Networks: The Price of Asynchrony*, UPMC, August 2009, <http://arxiv.org/abs/0908.0390>, Technical report.
- [57] Z. BOUZID, M. G. POTOP-BUTUCARU, S. TIXEUIL. *Optimal Byzantine Resilient Convergence in Asynchronous Robot Networks*, arXiv, June 2009, <http://arxiv.org/abs/0906.0651>, Technical report.
- [58] Z. BOUZID, M. G. POTOP-BUTUCARU, S. TIXEUIL. *Optimal byzantine resilient convergence in oblivious robot networks*, arXiv, May 2009, <http://arxiv.org/abs/0905.3967>, Technical report.
- [59] C. COTI, T. HÉRAULT, D. GROEN, M. MAMONSKI. *D1.2c: Adapted Verison of the OpenMPI Communication Library*, European Union, QosCosGrid project FP6-IST-2005-033883, June 2009, Technical report.
- [60] S. DEVISMES, F. PETIT, S. TIXEUIL. *Optimal Probabilistic Ring Exploration by Asynchronous Oblivious Robots*, n^o inria-00360305, INRIA, February 2009, <http://hal.inria.fr/inria-00360305/fr/>, Technical report.
- [61] S. DONFACK, L. GRIGORI, A. K. GUPTA. *Implementing communication-avoiding LU and QR factorizations for tall and skinny matrices on multicore architectures*, INRIA, 2009, Technical report.
- [62] S. DUBOIS, M. G. POTOP-BUTUCARU, S. TIXEUIL. *Dynamic FTSS in Asynchronous Systems: the Case of Unison*, n^o arXiv:0904.4615 (<http://arxiv.org/pdf/0904.4615v1>), INRIA, April 2009, <http://arxiv.org/pdf/0904.4615v1>, Research Report.
- [63] P. KUMAR, L. GRIGORI, F. NATAF, Q. NIU. *Combinative preconditioning based on Relaxed Nested Factorization and Tangential Filtering preconditioner*, n^o TR 6955, INRIA, June 2009, <http://hal.archives-ouvertes.fr/inria-00392881/en/>, submitted to Applied Numerical Mathematics Journal, Technical report.
- [64] M. NESTERENKO, S. TIXEUIL. *Ideal Stabilization*, arXiv, June 2009, <http://hal.inria.fr/inria-00394118/en/>, Research Report.

Other Publications

- [65] T. MASUZAWA, S. TIXEUIL. *Stabilizing Maximal Independent Set in Unidirectional Networks is Hard*, 2009, <http://arxiv.org/abs/0903.3106v1>.

References in notes

- [66] *EDonkey Homepage*, http://mldonkey.sourceforge.net/Main_Page, http://mldonkey.sourceforge.net/Main_Page.
- [67] K. AIDA, A. TAKEFUSA, H. NAKADA, S. MATSUOKA, S. SEKIGUCHI, U. NAGASHIMA. *Performance evaluation model for scheduling in a global computing system*, in "International Journal of High Performance Computing Applications", vol. 14, No. 3, 2000, p. 268-279, <http://dx.doi.org/10.1177/109434200001400308>.

- [68] A. D. ALEXANDROV, M. IBEL, K. E. SCHAUSER, C. J. SCHEIMAN. *SuperWeb: Research Issues in JavaBased Global Computing*, in "Concurrency: Practice and Experience", vol. 9, n^o 6, June 1997, p. 535–553.
- [69] L. ALVISI, K. MARZULLO. *Message Logging: Pessimistic, Optimistic and Causal*, 2001, Proc. 15th Int'l Conf. on Distributed Computing.
- [70] D. P. ANDERSON. *BOINC*, <http://boinc.berkeley.edu/>, <http://boinc.berkeley.edu/>.
- [71] A. BARAK, O. LA'ADAN. *The MOSIX multicomputer operating system for high performance cluster computing*, in "Future Generation Computer Systems", vol. 13, n^o 4–5, 1998, p. 361–372.
- [72] A. BARATLOO, M. KARAU, Z. M. KEDEM, P. WYCKOFF. *Charlotte: Metacomputing on the Web*, in "Proceedings of the 9th International Conference on Parallel and Distributed Computing Systems (PDCS-96)", 1996.
- [73] J. BEAUQUIER, C. GENOLINI, S. KUTTEN. *Optimal reactive k-stabilization: the case of mutual exclusion*. In *Proceedings of the 18th Annual ACM Symposium on Principles of Distributed Computing*, may 1999.
- [74] J. BEAUQUIER, T. HÉRAULT. *Fault-Local Stabilization: the Shortest Path Tree.*, October 2002, Proceedings of the 21th Symposium of Reliable Distributed Systems.
- [75] G. BOSILCA, A. BOUTEILLER, F. CAPPELLO, S. DJILALI, G. FEDAK, C. GERMAIN, T. HÉRAULT, P. LEMARINIER, O. LODYGENSKY, F. MAGNIETTE, V. NÉRI, A. SELIKHOV. *MPICH-V: Toward a Scalable Fault Tolerant MPI for Volatile Nodes*, 2002, in IEEE/ACM SC 2002.
- [76] A. BOUTEILLER, F. CAPPELLO, T. HÉRAULT, G. KRAWEZIK, P. LEMARINIER, F. MAGNIETTE. *MPICH-V2: a Fault Tolerant MPI for Volatile Nodes based on Pessimistic Sender Based Message Logging*, November 2003, in IEEE/ACM SC 2003.
- [77] A. BOUTEILLER, P. LEMARINIER, G. KRAWEZIK, F. CAPPELLO. *Coordinated Checkpoint versus Message Log for fault tolerant MPI*, December 2003, in IEEE Cluster.
- [78] T. BRECHT, H. SANDHU, M. SHAN, J. TALBOT. *ParaWeb: Towards World-Wide Supercomputing*, in "Proceedings of the Seventh ACM SIGOPS European Workshop on System Support for Worldwide Applications", 1996.
- [79] R. BUYYA, M. MURSHED. *GridSim: A Toolkit for the Modeling and Simulation of Distributed Resource Management and Scheduling for Grid Computing*, Wiley Press, May 2002.
- [80] J. W. BYERS, M. LUBY, M. MITZENMACHER, A. REGE. *A Digital-Fountain Approach to Reliable Distribution of Bulk Data*, in "proc. of the ACL SIGCOMM", 1998.
- [81] COSM. *Mithral Communications & Design Inc.*, <http://www.mithral.com/>, <http://www.mithral.com/>.
- [82] N. CAMIEL, S. LONDON, N. NISAN, O. REGEV. *The POPCORN Project: Distributed Computation over the Internet in Java*, in "Proceedings of the 6th International World Wide Web Conference", April 1997.

-
- [83] J. CAO, STEPHEN A. JARVIS, S. SAINI, GRAHAM R. NUDD. *GridFlow: Workflow Management for Grid Computing*, in "Proceedings of the Third IEEE/ACM International Symposium on Cluster Computing and the Grid", May 2003.
- [84] H. CASANOVA. *Simgrid: A Toolkit for the Simulation of Application Scheduling*. In *Proceedings of the IEEE International Symposium on Cluster Computing and the Grid (CCGrid '01)*, May 2001.
- [85] H. CASANOVA, A. LEGRAND, D. ZAGORODNOV, F. BERMAN. *Heuristics for Scheduling Parameter Sweep Applications in Grid Environments*, in "Proceedings of the Ninth Heterogeneous Computing Workshop", IEEE Computer Society Press, 2000, p. 349-363.
- [86] K. M. CHANDY, L. LAMPORT. *Distributed Snapshots: Determining Global States of Distr. systems.*, 1985, ACM Trans. on Comp. Systems, 3(1):63–75.
- [87] Y. CHEN, J. EDLER, A. GOLDBERG, A. GOTTLIEB, S. SOBTI, P. YIANILOS. *A prototype implementation of archival intermemory*, August 1999, In Proceedings of ACM Digital Libraries..
- [88] L. CHOY, S. G. PETITON, M. SATO. *Resolution of large symmetric eigenproblems on a world wide grid*, in "Seventh IEEE International Symposium on Cluster Computing and the Grid (CCGrid 2007), Rio de Janeiro, Brazil", IEEE Computer Society, May 2007, p. 301-308.
- [89] L. CHOY, S. G. PETITON, M. SATO. *Toward power-aware computing with dynamic voltage scaling for heterogeneous platforms*, in "Sixth International Workshop on Algorithms, Models and Tools for Parallel Computing on Heterogeneous Networks (HeteroPar) in conjunction with the 2007 IEEE International Conference on Cluster Computing (Cluster07), Austin, Texas USA", IEEE Computer Society Press, September 2007.
- [90] B. O. CHRISTIANSEN, P. CAPPELLO, M. F. IONESCU, M. O. NEARY, K. E. SCHAUSER, D. WU. *Javelin: Internet-Based Parallel Computing Using Java*, in "Concurrency: Practice and Experience", vol. 9, n^o 11, November 1997, p. 1139–1160.
- [91] B. COHEN. *Incentives Build Robustness in BitTorrent*, in "Workshop on Economics of Peer-to-Peer Systems, Berkeley", 2003.
- [92] S. DOLEV. *Self-stabilization*, 2000, M.I.T. Press.
- [93] G. FEDAK, C. GERMAIN, V. NÉRI, F. CAPPELLO. *XtremWeb: A Generic Global Computing System*, in "CCGRID'01: Proceedings of the 1st International Symposium on Cluster Computing and the Grid", IEEE Computer Society, 2001, 582.
- [94] I. FOSTER, A. IAMNITCHI. *On Death, Taxes, and the Convergence of Peer-to-Peer and Grid Computing*, in "2nd International Workshop on Peer-to-Peer Systems (IPTPS'03), Berkeley, CA", February 2003.
- [95] V. K. GARG. *Principles of distributed computing*, John Wiley and Sons, May 2002.
- [96] C. GENOLINI, S. TIXEUIL. *A lower bound on k-stabilization in asynchronous systems*, October 2002, Proceedings of the 21th Symposium of Reliable Distributed Systems.

-
- [97] D. P. GHORMLEY, D. PETROU, S. H. RODRIGUES, A. M. VAHDAT, T. E. ANDERSON. *GLUnix: A Global Layer Unix for a Network of Workstations*, in "Software Practice and Experience", vol. 28, n^o 9, 1998, p. 929–961.
- [98] C. GKANTSIDIS, P. RODRIGUEZ. *Network Coding for Large Scale Content Distribution*, in "Proceedings of IEEE/INFOCOM 2005, Miami, USA", March 2005.
- [99] L. GRIGORI, J. DEMMEL, X. LI. *Parallel Symbolic Factorization for Sparse LU Factorization with Static Pivoting*, in "SIAM Journal on Scientific Computing", vol. 29, n^o 3, 2007, p. 1289-1314.
- [100] B. HUDZIA. *Use of Multicast in P2P Network thought Integration in MPICH-V2*, Master of Science Internship, Pierre et Marie Curie University, September 2003, Technical report.
- [101] D. E. KEYES. *A Science-based Case for Large Scale Simulation, Vol. 1, Office of Science, US Department of Energy, Report Editor-in-Chief*, July 30 2003.
- [102] J. D. KUBIATOWICZ, D. BINDEL, Y. CHEN, P. EATON, D. GEELS, R. GUMMADI, S. RHEA, H. WEATHERSPOON, W. WEIMER, C. WELLS, B. ZHAO. *OceanStore: An Architecture for Global-scale Persistent Storage*, in "Proceedings of ACM ASPLOS", ACM, November 2000.
- [103] S. KUTTEN, B. PATT-SHAMIR. *Stabilizing time-adaptive protocols. Theoretical Computer Science 220(1)*, 1999.
- [104] S. KUTTEN, D. PELEG. *Fault-local distributed mending. Journal of Algorithms 30(1)*, 1999.
- [105] N. LEIBOWITZ, M. RIPEANU, A. WIERZBICKI. *Deconstructing the Kazaa Network*, in "Proceedings of the 3rd IEEE Workshop on Internet Applications WIAPP'03, Santa Clara, CA", 2003.
- [106] M. LITZKOW, M. LIVNY, M. MUTKA. *Condor — A Hunter of Idle Workstations*, in "Proceedings of the Eighth Conference on Distributed Computing, San Jose", 1988.
- [107] NANCY A. LYNCH. , M. KAUFMANN (editor)*Distributed Algorithms*, 1996.
- [108] MESSAGE PASSING INTERFACE FORUM. *MPI: A message passing interface standard*, June 12 1995, Technical report, University of Tennessee, Knoxville.
- [109] N. MINAR, R. MURKHART, C. LANGTON, M. ASKENAZI. *The Swarm Simulation System: A Toolkit for Building Multi-Agent Simulations*, 1996.
- [110] H. PEDROSO, L. M. SILVA, J. G. SILVA. *Web-Based Metacomputing with JET*, in "Proceedings of the ACM", 1997.
- [111] S. G. PETITON, L. CHOY. *Eigenvalue Grid and Cluster Computations, Using Task Farming Computing Paradigm and Data Persistency*, in "SIAM conference on Computational Science & Engineering (CSE'07), Costa Mesa, California, USA", February 2007.
- [112] PLATFORM. *Platform Computing - Accelerating Intelligence - Grid Computing*, <http://www.platform.com>, <http://www.platform.com/>.

- [113] D. QIU, R. SRIKANT. *Modeling and Performance analysis of BitTorrent-like Peer-to-Peer Networks*, in "SIGCOMM Comput. Commun. Rev.", vol. 34, n^o 4, 2004, p. 367–378.
- [114] B. QUÉTIER, M. JAN, F. CAPPELLO. *One step further in large-scale evaluations: the V-DS environment*, n^o RR-6365, INRIA, December 2007, <http://hal.inria.fr/inria-00189670>, Research Report.
- [115] S. RATNASAMY, P. FRANCIS, M. HANDLEY, R. KARP, S. SHENKER. *A Scalable Content Addressable Network*, in "Proceedings of ACM SIGCOMM 2001", 2001.
- [116] A. L. ROSENBERG. *Guidelines for Data-Parallel Cycle-Stealing in Networks of Workstations I: On Maximizing Expected Output*, in "Journal of Parallel Distributed Computing", vol. 59, n^o 1, 1999, p. 31–53.
- [117] A. ROWSTRON, P. DRUSCHEL. *Pastry: Scalable, Decentralized Object Location, and Routing for Large-Scale Peer-to-Peer Systems*, in "IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)", 2001, p. 329–350.
- [118] L. F. G. SARMENTA, S. HIRANO. *Bayanihan: building and studying Web-based volunteer computing systems using Java*, in "Future Generation Computer Systems", vol. 15, n^o 5–6, 1999, p. 675–686.
- [119] S. SAROIU, P. K. GUMMADI, S. D. GRIBBLE. *A Measurement Study of Peer-to-Peer File Sharing Systems*, in "Proceedings of Multimedia Computing and Networking, San Jose, CA, USA", January 2002.
- [120] SCIDAC. *SciDAC*, <http://www.scidac.org>, <http://www.scidac.org/>.
- [121] R. SHERWOOD, R. BRAUD, B. BHATTACHARJEE. *Slurpie: A Cooperative Bulk Data Transfer Protocol*, in "Proceedings of IEEE INFOCOM", March 2004.
- [122] J. F. SHOCH, J. A. HUPP. *The Worm Programs: Early Experiences with Distributed Systems*, in "Communications of the Association for Computing Machinery", vol. 25, n^o 3, March 1982.
- [123] O. SIEVERT, H. CASANOVA. *Policies for Swapping MPI Processes*, 2003, HPDC 2003, p.104–113.
- [124] I. STOICA, R. MORRIS, D. KARGER, F. KAASHOEK, H. BALAKRISHNAN. *Chord: A Scalable Peer-To-Peer Lookup Service for Internet Applications*, in "Proceedings of the 2001 ACM SIGCOMM Conference", 2001, p. 149–160.
- [125] G. TEL. *Introduction to distributed algorithms*, 2000, Cambridge University Press.
- [126] TERAGRID. *TeraGrid*, <http://www.teragrid.org>, <http://www.teragrid.org>.
- [127] B. UK, M. TAUFER, T. STRICKER, G. SETTANNI, A. CAVALLI. *Implementation and Characterization of Protein Folding on a Desktop Computational Grid - Is Charmm a Suitable Candidate for the United Devices Metaprocessor*, n^o 385, ETH Zurich, Institute for Computersystems, October 2002, Technical report.
- [128] Y.-M. WANG, W. K. FUCHS. *Optimistic Message Logging for Independent Checkpointing in Message-Passing Systems*, 1992, Symposium on Reliable Distributed Systems.

- [129] Y. YI, T. PARK, H. Y. YEOM. *A Causal Logging Scheme for Lazy Release Consistent Distributed Shared Memory Systems*, December 1998, In Proc. of the 1998 Int'l Conf. on Parallel and Distributed Systems.
- [130] Y. ZHANG, G. BERGÈRE, S. G. PETITON. *A parallel hybrid method of GMRES on Grid System*, in "Workshop on High Performance Grid Computing (HPGC'07), jointly published with IPDPS'07 proceedings, Long Beach, California, USA", March 2007.
- [131] B. Y. ZHAO, J. D. KUBIATOWICZ, A. D. JOSEPH. *Tapestry: An Infrastructure for Fault-tolerant Wide-area Location and Routing*, n° UCB/CSD-01-1141, UC Berkeley, April 2001, Technical report.
- [132] DATASYNAPSE. *Application Virtualization - DataSynapse Inc.*, <http://www.datasynapse.com>, <http://www.datasynapse.com/>.
- [133] GRIDSYSTEMS. *GRIDSYSTEMS - The Open Fabric of Virtualization*, <http://www.gridsystems.com>, <http://www.gridsystems.com/>.