

Hybrid Preemptive Scheduling of MPI Applications on Grids

Aurélien Bouteiller, Hinde-Lilia Bouziane, Thomas Herault, Pierre Lemarinier, Franck Cappello
INRIA/LRI, Université Paris-Sud, Orsay, France

E-mail: {bouteill, bouziane, herault, lemarini, fci}@lri.fr
phone: (+33) 1 69 15 4222
fax: (+33) 1 69 15 4213

Abstract—Time sharing between cluster resources in a Grid is a major issue in cluster and Grid integration. Classical Grid architecture involves a higher level scheduler which submits non-overlapping jobs to the independent batch schedulers of each cluster of the Grid. The sequentiality induced by this approach does not fit with the expected number of users and job heterogeneity of Grids. Time sharing techniques address this issue by allowing simultaneous executions of many applications on the same resources.

Co-scheduling and gang scheduling are the two best known techniques for time sharing cluster resources. Co-scheduling relies on the operating system of each node to schedule the processes of every application. Gang scheduling ensures that the same application is scheduled on all nodes simultaneously. Previous work has proven that co-scheduling techniques outperforms gang scheduling when physical memory is not exhausted.

In this paper, we introduce a new hybrid sharing technique providing checkpoint-based explicit memory management. It consists in co-scheduling parallel applications within a set, until the memory capacity of the node is reached, and using gang scheduling related techniques to switch from one set to another one. We compare experimentally the merits of the three solutions: Co, Gang and Hybrid Scheduling, in the context of out-of-core computing, which is likely to occur in the Grid context, where many users share the same resources. Additionally, we address the problem of heterogeneous applications by comparing hybrid scheduling to an optimized version relying on paired-scheduling. The experiments show that the hybrid solution is as efficient as the co-scheduling technique when the physical memory is not exhausted, can benefit from paired-scheduling optimization technique when applications are heterogeneous, and is more efficient than gang scheduling and co-scheduling when physical memory is exhausted.

Index Terms—Scheduling, Performance, MPI, Time sharing, Grid, Gang scheduling, Co-scheduling

I. INTRODUCTION

Two of the most fundamental principles of Grids are 1) the capacity to establish virtual organizations spanning over several administration domains in order to extend the number of resources accessible by users and 2) the coordination of these resources in order to cooperate in solving user problems. From these two principles eventually arises the need for efficient and fair resource sharing mechanisms.

The first principle inevitably tends to increase the pressure on the system mechanisms implementing the resource sharing between users. In the general situation, the resources belong to institutions and are already used by some of their members. Thus, building virtual organizations on top of already used resources increases the number of potential users for these resources, leading to an increased requirement to share these resources fairly among the users. In large systems within HPC centers, queues of jobs are already quite long. It is not uncommon to wait days before having large jobs done. If nothing is done in Grids associating tens of such sites, the waiting time would certainly evolve from days to weeks which in some circumstances will not be acceptable for users. Another fairness issue concerns the capacity to establish several level of priority among parallel jobs on the Grid. High priority jobs should be able to preempt the resource of a low priority parallel application under execution.

The second principle, as examined in the light of the first one, implies the use of efficient coordination mechanisms ensuring A) parallel application performance close to the one obtained in a dedicated system and B) limiting strictly the waste of computing resources by providing rapid parallel applications context switch. In this paper, we restrict our investigation domain to users running MPI applications on clusters shared within a Grid virtual organization. We focus on the following scenario: users submit their MPI jobs to a meta-scheduler (from a portal) which schedules them on dynamically selected clusters.

Any MPI execution may span over several clusters or stay within a single cluster. In the first case, an efficient coordination mechanism should schedule simultaneously all the components involved in each MPI execution. The efficiency in fulfilling criteria A) and B) depends on the synchronization mechanism coordinating the scheduling of the MPI subparts over all the clusters involved in the execution and on the speed of context switch between the MPI executions on each cluster. In the second case, clusters are not synchronized and the efficiency of the Grid only depends on the capability of each cluster management system to meet the criteria A) and B).

Fairness and performance are in principle contradictory. Reaching top performance on parallel execution involves dedicated usage of the cluster resources. The less the operating system interrupts the application execution, the higher is the performance. Usually, fairness relies on context switch mechanisms enabling resource sharing between users. Context switching adds an overhead on the total application execution time. The more context switches are experienced during an execution, the more the application is slowed down.

Sharing fairly the cluster resources between multiple users may be examined in two cases: 1) when the concurrent executions of all users fit in the memory of the cluster nodes and 2) conversely when disk storage should be used in addition to the memory in the cluster nodes to store all concurrent executions. We will call these two cases respectively in-core and out-of-core context switch.

The two principles lead to three main consequences: 1) a fast mechanism for switching the context of MPI execution on a cluster is the cornerstone to meet efficiency criteria, 2) because fairness and performance are contradictory objectives, we should consider a metric representing a tradeoff between them. For sake of simplicity, in this paper, we will consider application with the same execution time, and 3) In a Grid with lot of users, it is likely that out-of-core context switching will be the general case. Thus in this paper, we will essentially focus on this context.

In this paper we study several MPI application context-switching techniques, trying to discover which one has the lowest impact on application performance. We will demonstrate that the best technique for out-of-core context switching is a hybrid (two level) one mixing checkpoint based context switching of sets of MPI executions and uncoordinated scheduling (co-scheduling) of MPI executions within each set. We also demonstrate that Hybrid scheduling can benefit from the optimizations to gang and co-scheduling presented in previous works (paired scheduling).

The second section presents the related works. Section

III presents the different scheduling approaches compared in this paper. Section IV presents the general framework used to compare the different scheduling techniques. Section V presents the experimental results and section VI concludes and sums up what we learned from the experiences.

II. RELATED WORK

There are several main techniques to implement parallel application context switch in practice. One of the most used is batch scheduling, queuing the jobs submitted by the different users. In the general case, after being elected for execution, parallel jobs are scheduled sequentially (one after the other). Thus, only one parallel execution runs on the cluster at a given time. Examples of existing implementations of batch schedulers are PBS[1], LSF, Condor, etc. The main drawback of the batch scheduling approach is its lack of fairness between users submitting heterogeneous jobs. An application that needs a large number of nodes but for a short period may have to wait until all longer jobs running on fewer number of nodes terminate before being allowed to run.

This paper focuses on another family of approaches which lets the operating system schedule the processes of several parallel executions launched concurrently, according to their priorities. These techniques are called gang scheduling and co-scheduling, depending on the scheduling coordination of parallel execution processes. There is no coordination in co-scheduling. In gang scheduling, processes of a given application are scheduled simultaneously, requiring some synchronization mechanism. In these techniques all concurrent parallel applications reside in the cluster memory until their completion, generally leading to a huge demand on the virtual memory system.

[2], [3] has proposed one of the first implementation of gang scheduling, called SCore. SCore targets clusters and is based on a *Network preemption* procedure relying on the PM communication library [4]. The gang scheduling itself is performed using UNIX signal mechanism. When an application has to be unscheduled, all its processes on all nodes receive a SIGSTOP signal. Thus when another application is scheduled by the reception of the SIGCONT signal, it gets exclusive usage of computational power and network resources. However, the memory is shared between running and stopped applications. Memory sharing is resolved by the virtual memory mechanism of the operating system, as inactive applications may be transferred in swap memory. The gang scheduling strategy we present in this paper explicitly stores and reloads stopped processes

and does not rely on operating system swapping mechanism. In our checkpoint-based scheduling, typically only one execution resides in memory at a given time, thus limiting the memory sharing between applications. Another difference compared to SCore is the network flush algorithm. SCore does not use the Chandy & Lamport algorithm to flush the network, but a three phases synchronization algorithm using the PM flow-control protocol. The Chandy & Lamport algorithm we implemented uses only one synchronization phase. Note that uploading and downloading executions to and from the memory involves disk operations which can add a significant overhead.

An example of gang scheduling evaluation on LLNL's Cray T3D can be found in [5].

[6] evaluates different scheduling policies using the LSF batch scheduler. All policies are refinements of gang scheduling techniques allowing each application to run solely on the required processors. Three classes of applications are considered: short (5 minutes termination expected time), medium (60 minutes) and long running (no limit). The paper studies different policies when an application with a shorter termination expected time is queued. They demonstrate that preemption, implemented with checkpointing techniques, is crucial to obtain good response time. In this paper we extend the notion of hybridness, mixing gang scheduling with co-scheduling. The resulting approach could be adapted to work within a batch scheduler. The experimental studies concerning this issue will be presented in a future paper.

Co-scheduling involves launching all applications on the system resources and letting each node's operating system schedule the different jobs. The lack of coordination for the simultaneous execution of each node's part of an application is a major drawback for synchronous operations, but this approach allows a better overlapping of the communication of one job with the computation of another one (we assume that the communications are buffered independently of the process scheduling, as in kernel level protocol stacks). As co-scheduling relies on the operating system's memory scheduling, running out-of-core applications leads to high overhead due to the system swap policy. [7] introduces different paging techniques to reduce the number of page faults.

Some studies have compared and mixed gang scheduling and co-scheduling. In [8] the authors demonstrate experimentally that co-scheduling outperforms gang scheduling for clusters and in-core running applications. Our study pushes this result in the case of out-of-core applications, a major concern for Grid systems as the number of users and tasks expands. [9], [10], [11] consider specific classes of application based on

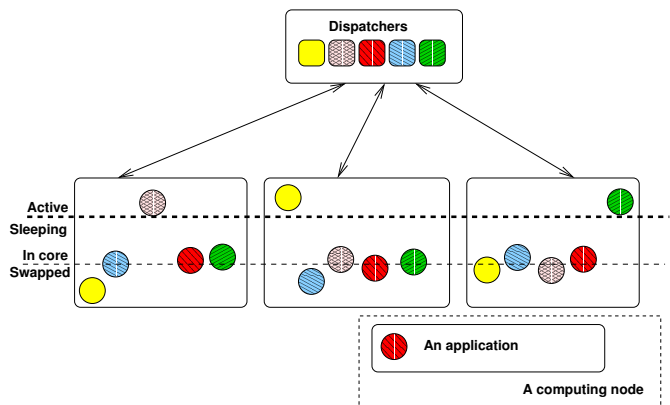


Fig. 1. Typical deployment of co-scheduling

their communication and I/O characteristics. Thus they improve gang scheduling of such known applications by co-scheduling applications of different classes. In this paper, we demonstrate that mixing the two techniques provides the best performance in the out-of-core case, even if the applications are identical.

III. THREE TECHNIQUES OF SCHEDULING

We study three different kinds of scheduling for sharing multiple MPI applications on a single set of nodes. Two of them, co-scheduling and gang scheduling, were compared in the context of clusters and using applications with low memory usage in [8]. These two techniques were not studied in the context of a large memory usage, leading to out-of-core computation when many applications runs simultaneously. We propose a new hybrid method based on the mix of the co-scheduling and the gang scheduling. The main idea is to limit the number of concurrent co-scheduled applications using gang scheduling global time slice and checkpoint related techniques, so that physical memory would not be exhausted by co-scheduling all applications.

A. Co-scheduling

The first one, called co-scheduling, consists of an uncoordinated approach. Figure 1 presents an exemple of deployment of the co-scheduling architecture. In this exemple we suppose that three simultaneous applications fits in physical memory. The processes of each application are launched simultaneously on all nodes. On each node, the local operating system scheduler is in charge of sharing resources among the processes of different applications. Thus processes of an application may not be scheduled at the same time on all nodes, which could impact performances of applications using a tightly synchronized communication scheme. Moreover, the frequent context switches between processes may introduce many cache faults. Nonetheless, this technique

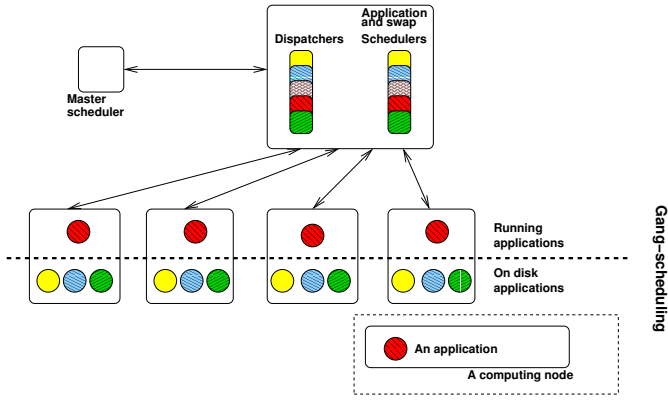


Fig. 2. Typical deployment of gang-scheduling

requires no specific implementation, and computational and network resources may be better used, as in the multi-threaded programming scheme.

B. Gang scheduling

The second approach is called gang scheduling. Figure 2 presents a typical deployment of the gang-scheduling architecture. It consists in synchronizing all the local schedulers so that all processes of a single distributed application are scheduled simultaneously, while all other applications are stopped and sleeping. All resources are thus employed on the execution of a single application, avoiding the wait for messages from a process not scheduled on another node. However, no multi-thread effect is possible as only one application is running at a time. However, there is no cache fault, and all physical memory is dedicated to the currently scheduled application. As discussed in [8], on many applications, the co-scheduling technique outperforms gang scheduling, as most applications do not perfectly overlap communications with computation.

Gang scheduling is implemented in the Master Scheduler of the MPICH-V framework.

C. New hybrid technique of time-sharing for high memory requirement

In this method, we propose to use co-scheduling for in-core computation, as it has been proved to be more efficient than gang scheduling. When out-of-core computation would appear using co-scheduling, we use a gang scheduling related technique to enforce that only a subset of the applications is running on the nodes. As the overhead induced by the applications switch is related to time to store process memory on local disk, it is much higher than node operating system context switch overhead. As a consequence, time slices have to be much longer. Thus, gang scheduling is mandatory to reach good performance for any communicating

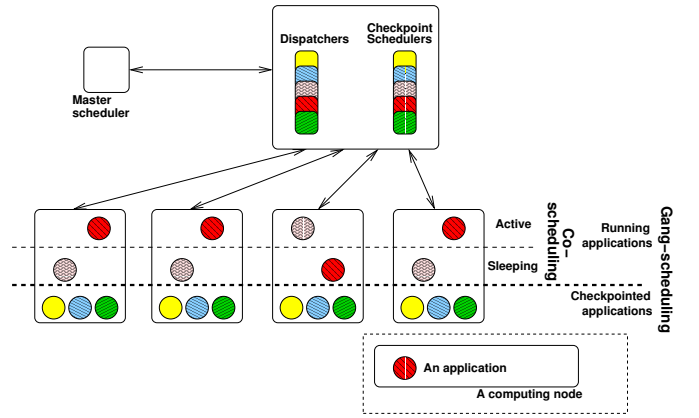


Fig. 3. Typical deployment of hybrid scheduling

application. Waiting for a message during the full time slice of an application would lead to very poor network performance.

Figure 3 presents a typical deployment of the hybrid architecture. In this example we suppose that physical memory size allows us to run two co-scheduled applications simultaneously without inducing out-of-core execution.

The Master Scheduler controls Dispatchers and Checkpoint Schedulers of each application, forcing some applications to be stopped and swapped out of memory, while some others are restarted. We explain in section III-C.2 different methods to swap an application out of memory. For the set of applications that are running, their processes are co-scheduled by each node's operating system.

1) Network management for application switching:

To stop an application, the network status has to be saved. We use the Chandy&Lampert algorithm to flush the network before the application is stopped. The Checkpoint Scheduler and the communication daemons are designed to save the network state. The checkpoint scheduler requests every communication daemon to take a global snapshot by sending them a tag. On the reception of a tag, a daemon stops the computing process, and then sends the tag in every communication channel (including checkpoint scheduler). When the checkpoint scheduler have received a tag from every communication daemon, the network flush is achieved.

2) Memory management for application switching:

a) *SIGSTOP/SIGCONT policy, system memory management:* In previous implementations of gang scheduling, each process of an application is stopped by the SIGSTOP UNIX signal. In the case of large memory consumption, memory sharing is managed by the operating system's swap policy. As running applications were swapped out to disk during the previous time slice, the memory pages are reloaded on demand during

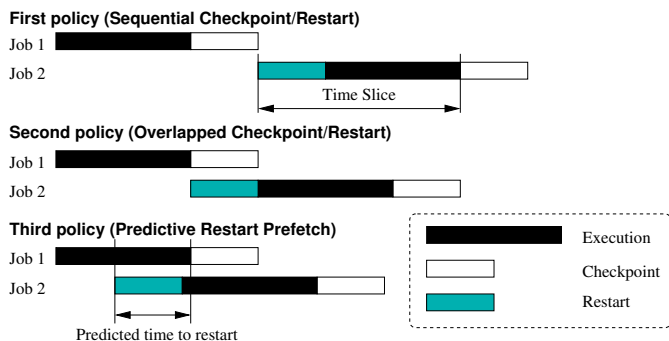


Fig. 4. Policy for checkpointing a job and restarting the next scheduled one on the same node

the execution (thus swapping out some pages used by stopped applications). The number of page faults has a major impact on overall performance. The overhead using this technique is very unpredictable as it relies on operating system swapping policy. Moreover, relying on SIGSTOP/SIGCONT mechanism would introduce perturbations in our page fault measurements, as it is difficult to differentiate out-of-core computation and context switch induced page faults. Thus, we prefer to use a checkpoint based technique, whose overhead is well bounded.

b) Checkpoint policies, explicit memory management: In our implementation, we use checkpoint/restart to explicitly manage memory swapping of applications memory. When a computing process is requested to be stopped, it performs a checkpoint to local disk, thus freeing all memory it uses. The complete memory of the process to be scheduled in the next time slice is reloaded from checkpoint, thus no page fault occurs during the time slice. Incremental checkpoint related techniques may be used to improve checkpoint performances. This memory management is equivalent to aggressive paging out technique described in [7].

Many policies may be used to overlap checkpoint, restart, and computations. Figure 4 presents the three techniques implemented in the comparison framework. The first method (called sequential checkpoint/restart) does not overlap checkpointing and restarting. It avoids loading the two applications simultaneously in memory at the expense of serializing checkpoint, restart and computation during the context switch.

The second technique (called overlapped checkpoint/restart) overlaps checkpoint and restart, trying to reduce context switch cost. It requires more memory as one application is reloaded before the previous is fully flushed out of memory, and induces simultaneous disk accesses.

The third technique (called predictive restart prefetch) tries to prefetch restart during the end of the time slice of the previous application, so that restart is overlapped

by computation of the application to be stopped, and checkpoint is overlapped by computation of the next application. It has to rely on an oracle, predicting time to restart, and induces simultaneous memory usage during the end of the time slice of the application to be stopped.

These three techniques are compared in section V-B.

3) *Paired hybrid scheduling:* In the context of pure gang-scheduling, [9] prove that scheduling two applications with very different communication/computation ratio during the same time slice improves performance greatly. This optimization can be used in the context of Hybrid scheduling of heterogeneous applications, by building the co-scheduled subset with consideration to the communication/computation ratio of the applications, and trying to avoid creating homogeneous subsets.

IV. COMMON FRAMEWORK

In this paper, we focus on comparing and mixing two scheduling methods for MPI application time sharing. The first one is based on the ability to stop and restart a set of MPI applications so that memory used by these applications is available to another set of applications. At the end of a time slice, a set of applications is dumped on disk, and another set of applications is loaded from a previous checkpoint and run during the next time slice. The second one is based on running all MPI applications simultaneously on the same set of nodes. This method relies on the operating system scheduler to perform time sharing and is called co-scheduling.

The MPICH-V framework offers both stop and restart capable and standard MPI implementations. It's main focus is comparison of different kind of fault tolerant protocols for MPI. Deviating from this main purpose, we developed two non fault tolerant protocols. One includes a distributed checkpoint facility based on the Chandy&Lamport algorithm (Vcl [12], [13]), and another one implements basic MPI communication without checkpoint capabilities (Vdummy). Using such an implementation is mandatory to perform a fair comparison: as the two implementations share the same framework, any performance difference is related to the scheduling itself, and not to implementation optimizations. Moreover, we compared the MPICH-V framework to the reference implementation MPICH-P4 in [13]. Figure 5 compares the architectures, while figure 7 shows the ping-pong performance comparison between the MPICH-V framework and the reference implementation, and validates the checkpoint enabled MPICH-Vcl version performance compared to standard MPICH-Vdummy version.

MPICH-V is based on the MPICH library [14], which builds a full MPI library from a channel. A channel implements the basic communication routines for a specific

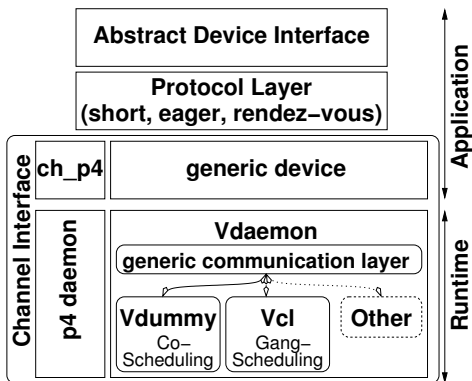


Fig. 5. Architecture of MPICH-V compared to architecture of MPICH-P4




Fig. 6. Typical deployments of MPICH-V with many applications

hardware or for new communication protocols. MPICH-V consists in a set of runtime components and a channel (ch_v) for the MPICH library.

The different protocols are implemented in the MPICH-V framework at the same level of the software hierarchy, between a MPI high level protocol management layer (managing global operations, point to point protocols, etc.) and the low level network transport layer. Among the other benefits, this allows us to keep unmodified the MPICH implementation of point to point and global operations, as well as complex concepts such as topologies and communication contexts. A potential drawback of this approach might be the necessity to implement a specific driver for all types of Network Interface (NIC). However, several NIC vendors provide low level, high performance (zero copy) generic socket interfaces such as Socket-GM for Myrinet, SCI-Socket for SCI and IPoIB for Infiniband. MPICH-V protocols typically sit on top of these low level drivers. So this is one of the most relevant layer for implementing new MPI capabilities if criteria such as design simplicity, high performance, heterogeneous network migration and portability are to be considered.

MPICH-V provides all the components necessary to stop and restart MPI applications. Some of them have been slightly modified to focus on the scheduling of MPI applications, while some have been added, specifically to manage sets of applications.

A. Dispatcher

The dispatcher of the MPICH-V environment has two main purposes: 1) to launch the whole runtime environment (encompassing the computing nodes and the auxiliary "special" nodes) on the pool of machines used for the execution, and 2) to monitor this execution, by detecting node disconnection, and then stop the execution.

The Dispatcher is in charge of a single MPI application. If more than one application is running at a time on a cluster, each is controlled by its own Dispatcher.

B. Driver

The driver is the part of the MPICH-V framework linked with the MPI application. It implements the Channel Interface of MPICH. Our implementation only provides synchronous functions (bsend, breceive, probe, initialize and finalize). The asynchronism of communication is delayed to the daemon.

C. Communication daemon

The core of the communication daemon is a select loop: it manages one socket for every computing node and one socket for every specific component. Every send or receive operation is asynchronous. Thus, a communication is not blocked by another slower one. On the contrary, the communication across the Inter Process Communication mechanism to the MPI process is synchronous and its granularity is the whole protocol message. The communication daemon is in charge of all distributed checkpoint mechanisms.

The checkpoint of the daemon uses an explicit serialization of its data and when a checkpoint is requested, some messages have to be logged on the receiver (considered as the in-transit messages of the Chandy-Lamport algorithm).

a) *The generic Communication daemon*: Daemons implement a generic communication layer to provide all the point-to-point communication routines between the different kind of components involved in the MPICH-V architecture, independently of the protocols. Checkpoint-enabled protocols are designed through the implementation of a set of hooks called in relevant routines of the generic layer and some specific components (figure 5).

The collection of all these functions is defined through a fault tolerance API and each protocol implements this API. In order to reduce the number of system calls, communications are packed using *iovec* related techniques by the generic communication layer. The different communication channels are multiplexed using a single thread and the *select()* system call. This common

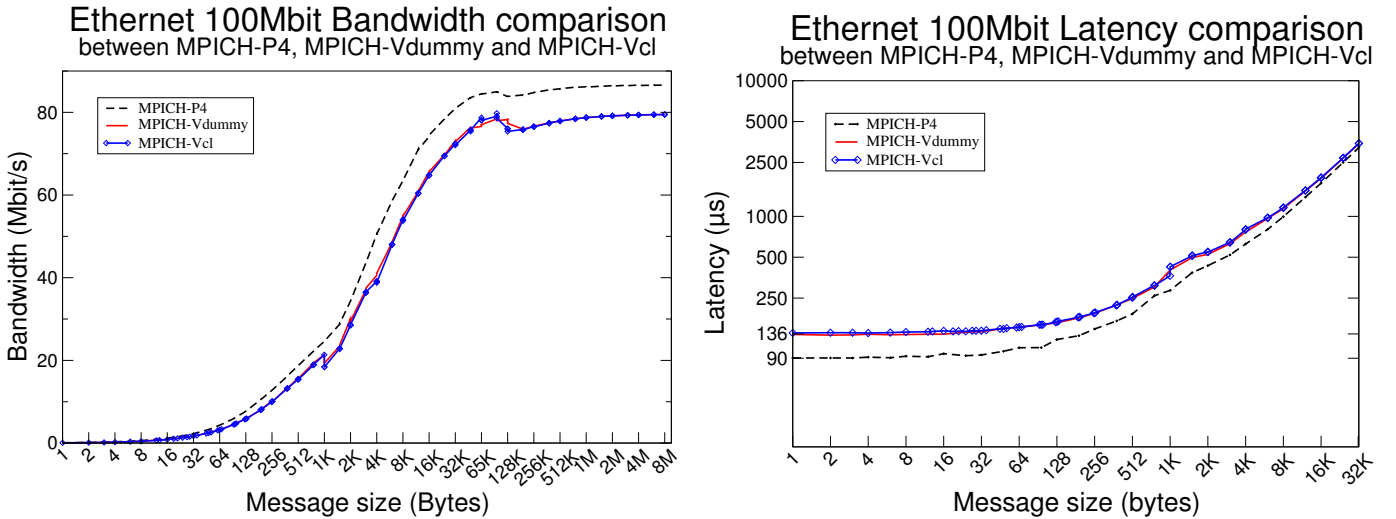


Fig. 7. Bandwidth and latency comparison between MPICH-P4, MPICH-Vdummy, and MPICH-Vcl on Fast-Ethernet network.

implementation of communications eases the implementation of protocols and allows a fair comparison between them.

D. Checkpoint scheduler

The checkpoint scheduler requests computation processes to checkpoint according to a given policy. The policy is protocol dependent. In this article, the checkpoint scheduler uses a coordinated checkpoint policy driven by the master scheduler, intended to checkpoint an application before it is stopped at the end of a time slice.

E. Master scheduler

The master scheduler is a new component introduced in the MPICH-V architecture in order to target Grid scheduling. The master scheduler coordinates the scheduling of all the applications on the system. Given a list of applications to schedule, it first launches the dispatcher and checkpoint scheduler components of each application. The application deployment itself is performed by the dispatcher of this application (figure 6).

The number n of job to run simultaneously is given as parameter of the master scheduler. The master scheduler associates a time slice to each application. When a scheduled application has been running long enough to expire its time slice, the master scheduler requests the dispatcher and the checkpoint scheduler to stop this application. When it is stopped, the master scheduler requests another dispatcher to restart the associated application, and the time slice for this application begins.

The master scheduler implements various policies for stopping/restarting applications. It is possible to Co-schedule k applications of a set of n . To perform set context switch, three checkpoint/restart overlap policies are

implemented. These policies are detailed in section III-C.2.

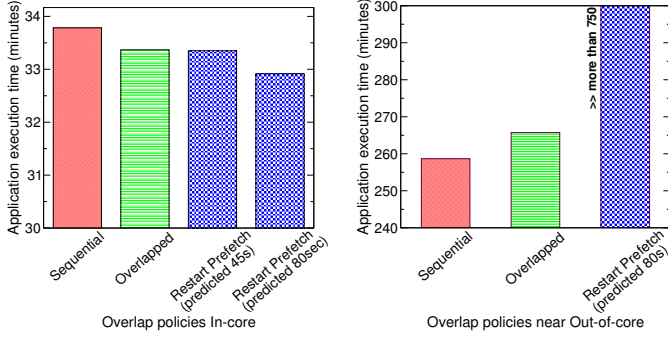
V. PERFORMANCE EVALUATION

A. Experimental conditions

Experiments are run on a 32-nodes cluster. Each node is equipped with an Athlon XP 2800+ processor, running at 2GHz, 1GB of main memory (DDR SDRAM), and a 70GB IDE ATA100 hard drive and a 100Mbit/s Ethernet Network Interface Card. Swap space is set to 10GB. All nodes are connected by a single Fast Ethernet Switch.

All these nodes use Linux 2.4.21 as operating system. The tests and benchmarks are compiled with GCC 2.95-5 (with flag `-O3`) and the PGI Fortran77 compilers. All tests are run in dedicated mode. Each measurement is repeated 5 times and we present a mean of them.

The first experiments are synthetic benchmarks analyzing the individual performance of the subcomponents. We use the NetPIPE [15] utility to measure bandwidth and latency of the MPICH-V framework. This is a ping-pong test for several message sizes and small perturbations around these sizes. The second set of experiments is the set of kernels and applications of the NAS Parallel Benchmark suite [16], written by the NASA NAS research center to test high performance parallel computers. These benchmarks cover a large panel of communication schemes: each benchmark tests a particular communication scheme, and communication computation ratio. CG benchmark presents heavy point-to-point latency driven communications; BT benchmark presents large point-to-point messages, and communications overlapped by computation. Each process of a BT benchmark uses 175MB of memory for class C on 25 nodes, 135MB for class B on 9 nodes, and each process of CG class C on 8 nodes uses 157MB. A node has 1GB



(a) 2 Hybrid scheduled BT.C.25, 1 running at a time (in-core).

(b) 10 Hybrid scheduled BT.C.25, 5 running at a time (near out-of-core).

Fig. 8. Checkpoint / Restart application context switch policies comparison, using NAS benchmark BT class C on 25 nodes on Ethernet performance criteria.

of memory (900 MB of user space memory). Thus up to five simultaneous applications fit in physical memory whatever the application we use. Moreover, with seven or more simultaneous applications, swapping is always needed.

As we compare scheduling algorithm in term of both performance and fairness, we propose to use the following tradeoff metric to measure performance: the slowdown between concurrent application set makespan and time to perform according number of sequential execution plus the standard deviation of individual concurrent execution time over the average concurrent execution time.

$$\frac{t_{\text{makespan}}}{n \cdot t_{\text{sequential}}} + \frac{\sqrt{\sum (t_{\text{concurrent}} - t_{\text{concurrent}})^2}}{n \cdot t_{\text{concurrent}}}$$

B. Checkpoint/restart overlap scheduling policy

Figure 8 presents the performance of NAS BT class C benchmark on 25 nodes for the three application context switch policies of Checkpoint/Restart presented in section III-C.2.b. On the first figure, two concurrent applications are gang scheduled. On the second figure, 5 applications of 10 are run simultaneously, thus filling physical memory.

When physical memory is not exhausted by computing applications, as expected, the more overlap is reached, the better performance is. Thus, the Overlapped Checkpoint/Restart method performs better than Sequential Checkpoint/Restart, and the Predictive Restart Prefetch method performs better than the two others. Moreover, it is better to use a greedy prefetch than an exact one, as the predicted checkpoint time may vary around the mean value.

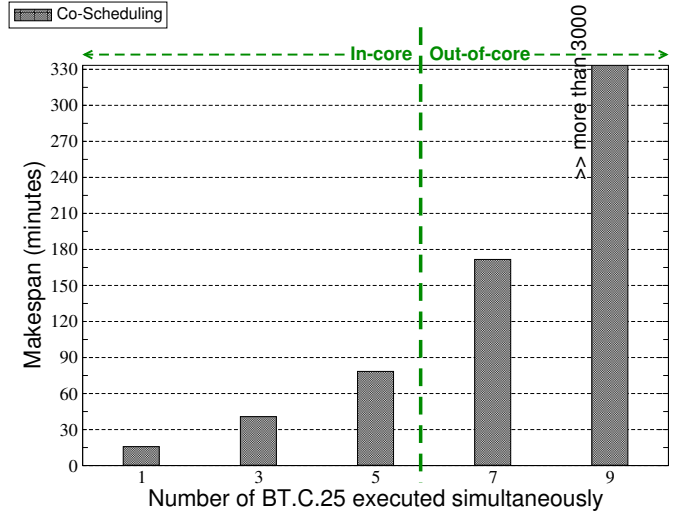


Fig. 9. Makespan for n simultaneous BT class C 25 nodes using Co-scheduling.

When running enough applications to exhaust physical memory, the overlapping strategies do not perform as well. The finishing application stays in memory during its checkpoint, thus inducing memory swapping for the restarted applications. Moreover, the restart prefetch strategy suffers from a dramatic overhead related to the very intensive memory usage induced by the simultaneous execution of applications at the end of their time slice and applications beginning a new time slice. Concurrent applications swapping on disk increase the checkpoint time, while checkpoint and restart accesses on disk decrease swapping performance. Thus, the more the disk is accessed simultaneously, the more the overall performance is decreased. As it sequentializes both disk access and memory usage, the sequential policy performs better when physical memory is near to be filled.

C. Scheduling techniques performance comparison

Figure 9 presents the computation time of n simultaneous BT class C on 25 nodes benchmark. When the memory used by the concurrent applications fits in-core, the computation throughput increases slightly with the number of applications. For 7 simultaneous applications, the overall throughput is decreasing, and for 9 simultaneous applications, the makespan is twenty times the sequential batch scheduling makespan. Figure 10 explains this result. It presents the average number of major page faults per minute for 7 and 9 simultaneous applications using Co-scheduling. On one hand, for 7 Co-scheduled applications, only a subset of the applications hits the swap: the first application does not suffer from any page fault (8.5 page faults per minutes), while another suffers from more than 1245 page faults per minutes. On the other hand, for 9 Co-scheduled applications, all applications hit the swap equally (Standard deviation is

Number of applications	Average number of major page faults for all nodes of each application, per minutes									Average page faults per minutes, All applications	Standard deviation
	app 0	app 1	app 2	app 3	app 4	app 5	app 6	app 7	app 8		
7	8.5	239.5	264.25	951	1145.25	1245.5	1074			704	474.9
9	484.58	405.94	564.78	524.4	510.66	577.26	506.94	481.84	509.4	507.4	47.1

Fig. 10. Page faults statistics for simultaneous BT class C on 25 nodes

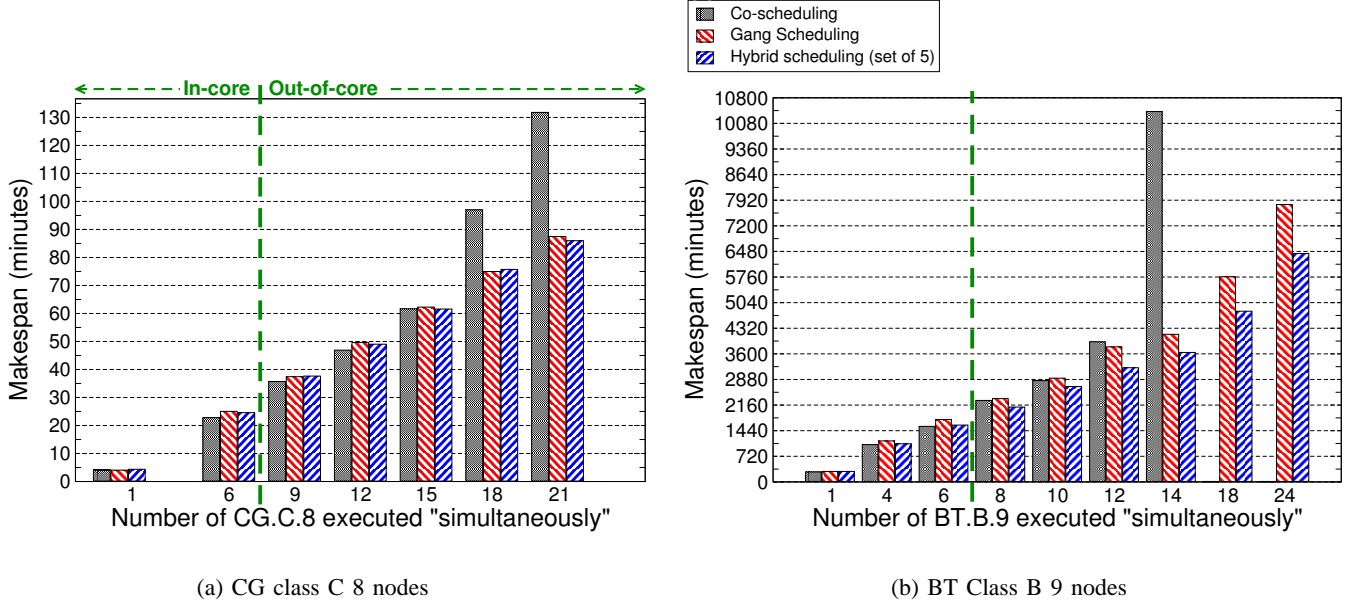


Fig. 11. Makespan for n simultaneous applications using Co-scheduling, Gang scheduling or hybrid scheduling.

less than 47.1 for an average page faults per minute of 507.4). This illustrates the lack of fairness of the virtual memory management used in the Linux 2.4.21 kernel when only a small amount of swap is used. This algorithm achieves good performance by scheduling more time for the in-core applications, at the cost of serializing executions. When memory occupation leads the size of the page cache table to reach its lower bound *pages_table_low*, the virtual memory manager applies a first "gentle" policy, which is the case for 7 simultaneous applications. For 9 simultaneous applications and above, the upper bound *pages_table_high* is reached, setting the virtual memory manager in an "aggressive" (but fair) swapping policy. In this case, all applications have fair access to physical memory, but performance suffers from a dramatic decrease. Even if these bounds may be tuned, at some point, the kernel has to switch to the aggressive policy, in order to ensure availability and fairness.

Figure 11 presents the computation time of n simultaneous NAS benchmarks, using co-scheduling, gang scheduling or hybrid scheduling. Hybrid scheduling relies on the operating system scheduler for executing a set of 5 applications simultaneously, in order to occupy all the physical memory without inducing out-of-core computation. The sequential Checkpoint/Restart policy is used to perform set of applications context switch. Time slices are 900 seconds.

For the CG benchmark (figure 11(a)), the co-

scheduling method reaches good performance up to use about 3 times the physical memory (2800MB). Comparing to the BT benchmark results (figure 9), this shows that the point of inefficiency of co-scheduling is tightly related to the pattern of memory accesses. The application often uses the same pages, reducing the number of page faults. However, for these two applications, the huge impact on co-scheduling of the kernel's aggressive swapping policy is observed. When co-scheduling is outperformed, gang and hybrid scheduling techniques reach the same performance. This shows that there is almost no benefit of co-scheduling a subset of applications for this benchmark. In CG, communication prevails, thus the bandwidth is divided between co-scheduled applications.

Figure 11(b) focuses on performance comparison between gang scheduling and hybrid scheduling when co-scheduling is outperformed due to out-of-core computation for the BT benchmark. For each time slice, gang scheduling and hybrid scheduling have to checkpoint and restart one application. During that time slice, no swap effect occurs, as the number of applications running simultaneously is either one for gang scheduling or however many fit in-core for hybrid scheduling. Obviously, the performance is exactly the performance of pure gang scheduling compared to pure in-core co-scheduling. Unlike the CG benchmark, computation prevails in the BT benchmark. For this kind of applications, hybrid

		co-sched	gang sched	hybrid sched
CG	in-core	0.92	1.02	0.99
	out-of-core	1.3	1.02	0.99
BT	in-core	0.9	1.02	0.90
	out-of-core	1.9	1.07	0.90

Fig. 12. Average slowdown using co, gang, or hybrid scheduling, compared to sequential batch scheduling

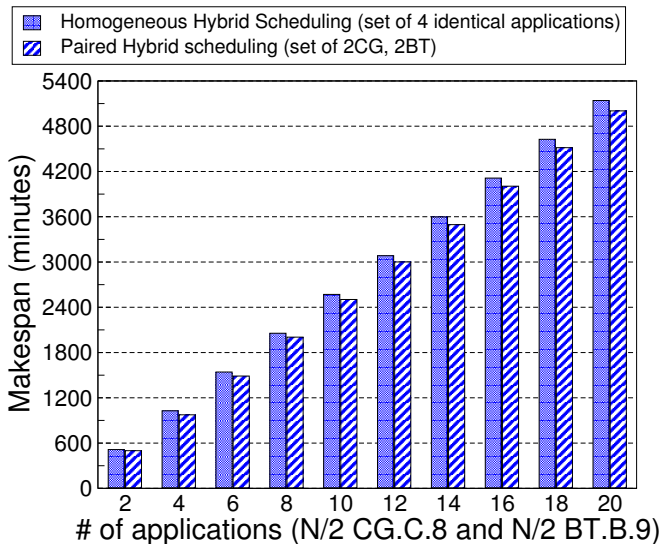


Fig. 13. Makespan for n simultaneous heterogeneous applications using homogeneous hybrid scheduling or paired hybrid scheduling.

scheduling performs 17% better than gang scheduling.

The figure 12 compares the relative slowdown of co, gang, and hybrid scheduling compared to ideal (overhead free) sequential batch scheduling. For the CG benchmark, gang and hybrid scheduling performs equally to sequential scheduling. Co-scheduling shows a 6 to 8% performance improvement up until we reach out-of-core computing, where it shows a 30% performance penalty. For the BT benchmark, in-core, co-scheduling and hybrid scheduling give a 10% improvement compared to sequential scheduling, while gang scheduling suffers from a 2% overhead. When out-of-core memory is used, co-scheduling performance is very poor, and gang-scheduling suffers from a 7% slowdown. In the same context, hybrid scheduling produce a 10% improvement compared to sequential scheduling. For every benchmark, the overall throughput of hybrid scheduling is equal or up to 10% better than sequential scheduling.

D. Case study for heterogeneous applications

Focusing on the case of heterogeneous sets of application, let consider a total of n applications, where $n/2$ are communication bound (CG.C.8 in this experiment) and $n/2$ are computation bound (BT.B.9). With regard to available physical memory, up to 5 of these applications may be co-scheduled during each time slice. However,

in order to have the same number of computing and communicating applications, we restricted the subset size to 4 in this experiment. On one hand, for homogeneous hybrid scheduling, the construction of each subset ensures that only 4 identical applications are running during each time slice. Thus half of the time slices are dedicated to only communicating applications, and half to computing applications. On the other hand, for paired hybrid scheduling, all subsets are identical and composed of 2 communicating and 2 computing applications. Figure 13 presents homogeneous hybrid scheduling compared to paired hybrid scheduling. Whatever is the total number of applications, considering the different communication/computation ratio of applications improves hybrid scheduling performance of 2.7%. Compared to the 17% performance improvement between gang and hybrid scheduling, this illustrates that most of the communication/computation overlap is due to co-scheduling, and only a small part is related to heterogeneousness.

VI. CONCLUSION

In this paper, we compare several scheduling techniques for the Grid. Some are well known, such as gang scheduling (which schedules a whole single application on all nodes at a given time) and co-scheduling (which schedules all applications simultaneously on all nodes). We propose a new scheduling approach, based on the two previous, that we call hybrid scheduling.

Hybrid scheduling splits the set of applications to schedule in subsets, co-scheduling the applications of a same subset and gang scheduling the different subsets. The main decisive factor is the amount of physical memory used by all the processes of a same subset. The technique uses user-mode checkpoints to stop and restart gang scheduled applications. Using three different policies, we studied experimentally the merits of overlapping or not the checkpoint and restart during gang scheduling context switch.

We conducted a set of experiments using the NAS parallel benchmarks. We show that for out-of-core computation, co-scheduling behaves accordingly to the swap policy of the operating system, and is eventually hit by very poor performance. We show that according to the application computation/communication ratio, hybrid scheduling compares favorably or equally to gang scheduling. We show that hybrid scheduling can benefit from scheduling simultaneously applications with different computation/communication ratios. Comparing to sequential batch scheduling, hybrid scheduling reaches 10% better or equal throughput depending on the kind of applications, and offers a better fairness.

We plan to improve the master dispatcher, in order to dynamically adapt the number of applications running concurrently to the memory occupation of nodes. Another issue is the gang scheduler context switch efficiency, so we plan to compare user-space checkpoint techniques to optimized kernel swap algorithm in the context of overflowed physical memory. Finally, we plan to integrate our hybrid scheduling system into a meta batch scheduling system for the grid, to compare the fairness, reactivity and performance of such a system with classical batch scheduling.

VII. ACKNOWLEDGMENTS

The MPICH-V project is founded by the French ministry of research through the "ACI Grid" initiative. We would like to thank Larry Carter, from the university of San Diego, who has kindly helped us to improve the overall quality of the paper by his thoughtful reading.

REFERENCES

- [1] R. L. Henderson, "Job scheduling under the portable batch system," in *Workshop on Job Scheduling Strategies for Parallel Processing*, 1995, pp. 279 – 294.
- [2] A. Hori, H. Tezuka, Y. Ishikawa, N. Soda, H. Konaka, and M. Maeda, "Implementation of gang-scheduling on workstation cluster," in *Proceeding of IPPS'96 Workshop on Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph, Eds. Springer-Verlag, April 1996, pp. 126–139.
- [3] A. Hori, H. Tezuka, and Y. Ishikawa, "Overhead analysis of preemptive gang scheduling," *Lecture Notes in Computer Science*, vol. 1459, pp. 217–230, April 1998.
- [4] H. Tezuka, A. Hori, Y. Ishikawa, and M. Sato, "PM: An operating system coordinated high performance communication library," in *Proceedings of the international conference and Exhibition on high-Performance Computing and Networking*. Springer-Verlag, April 1997, pp. 708–717.
- [5] D. G. Feitelson and M. A. Jette, "Improved utilization and responsiveness with gang scheduling," in *Job Scheduling Strategies for Parallel Processing*, D. G. Feitelson and L. Rudolph, Eds., vol. 1291. Springer Verlag, 1997, pp. 238–261.
- [6] E. W. Parsons and K. C. Sevcik, "Implementing multiprocessor scheduling disciplines," in *Proceeding of IPPS'97 Workshop on Job Scheduling Strategies for Parallel Processin*, D. G. Feitelson and L. Rudolph, Eds. Springer Verlag, April 1997, pp. 166–192.
- [7] K. D. Ryu, N. Pachapurkar, and L. L. Fong, "Adaptive memory paging for efficient gang scheduling of parallel applications," in *18th International Parallel and Distributed Processing Symposium IPDPS'04*, April 2004.
- [8] P. Strazdins and J. Uhlmann, "A comparison of local and gang scheduling on a beowulf cluster," in *Proceedings of the 2004 IEEE International Conference of Cluster Computing*. IEEE CS Press, September 2004, pp. 55–62.
- [9] Y. Wiseman and D. G. Feitelson, "Paired gang scheduling," in *IEEE Transactions on Parallel and Distributed Systems*, vol. 14, June 2003, pp. 581–592.
- [10] W. Lee, M. Frank, V. Lee, K. Mackenzie, and L. Rudolph, "Implications of I/O for gang scheduled workloads," *Proceedings of the Job Scheduling Strategies for Parallel Processing*, pp. 215–237, 1997.
- [11] F. A. B. da Silva and I. D. Scherson, "Improving parallel job scheduling using runtime measurements," *Proceedings of the 6th Workshop on Job Scheduling Strategies for Parallel Processing*, pp. 18–38, May 2000.
- [12] A. Bouteiller, P. Lemarinier, G. Krawezik, and F. Cappello, "Coordinated checkpoint versus message log for fault tolerant MPI," in *IEEE International Conference on Cluster Computing (Cluster 2003)*. IEEE CS Press, December 2003.
- [13] P. Lemarinier, A. Bouteiller, T. Herault, G. Krawezik, and F. Cappello, "Improved message logging versus improved coordinated checkpointing for fault tolerant MPI," in *IEEE International Conference on Cluster Computing (Cluster 2004)*. IEEE CS Press, 2004.
- [14] W. Gropp, E. Lusk, N. Doss, and A. Skjellum, "High-performance, portable implementation of the MPI message passing interface standard," *Parallel Computing*, vol. 22, no. 6, pp. 789–828, September 1996.
- [15] Q. Snell, A. Mikler, and J. Gustafson, "Netpipe: A network protocol independent performance evaluator," in *IASTED International Conference on Intelligent Information Management and Systems.*, June 1996.
- [16] D. Bailey, T. Harris, W. Saphir, R. V. D. Wijngaart, A. Woo, and M. Yarrow, "The NAS Parallel Benchmarks 2.0," Numerical Aerodynamic Simulation Facility, NASA Ames Research Center, Report NAS-95-020, 1995.

Aurélien Bouteiller is a PhD student in the Cluster and Grid group of the LRI laboratory at Paris-South university and is a member of the Grand-Large team of INRIA. He has obtained a Master in parallel computer science in 2002 from the french Paris-South university. He contributes to the MPICH-V project, a fault tolerant MPI implementation comparing different fault tolerant protocols. His research interests include high performance fault tolerant MPI and checkpoint optimizations and performance evaluation.

Hinde Lilia Bouziane has obtained a Master degree in computer science in September 2004 from the French Paris-South University. She spent five months in the Grand-Large research team of INRIA working on hybrid scheduling of MPI applications for clusters. In November 2004, she has started a PhD in the PARIS research team of INRIA, located at IRISA (Rennes, France). Her research interests concern advanced programming models for GRIDs, involving software component and workflow models.

Thomas Herault is associate professor at the Paris South University. He defended it's PhD on the mending of transient failure in self-stabilizing systems under the supervision of Joffroy Beauquier. He is a member of the Grand-Large INRIA team and works on fault-tolerant protocols in large scale distributed systems. He contributes to the MPICH-V project, whose goal is to evaluate and compare fault tolerant protocols for MPI.

Pierre Lemarinier has obtained a Master in computer science in 2002 from the french Paris-South University. He is a PhD student in the parallelism team and the Grand-Large team of the LRI laboratory of Paris-South. His research interests include fault tolerant protocols conception and validation and MPI implementations (MPICH-V).

Franck Cappello holds a Research Director position at INRIA, after having spent 8 years as CNRS researcher. He leads the Grand-Large project at INRIA and the Cluster and Grid group at LRI. He has authored more than 50 papers in the domains of High Performance Programming, Desktop Grids and Fault tolerant MPI. He is editorial board member of the "international Journal on GRID computing" and steering committee member of IEEE/ACM CCGRID. He organizes annually the Global and Peer-to-Peer Computing workshop. He has initiated and heads the XtremWeb (Desktop Grid) and MPICH-V (Fault tolerant MPI) projects. He is currently involved in two new projects: Grid eXplorer (a Grid Emulator) and Grid5000 (a Nation Wide Experimental Grid Testbed).