

Application Execution Management on the InteGrade Opportunistic Grid Middleware

Francisco José da Silva e Silva^a, Fabio Kon^b, Alfredo Goldman^b, Marcelo Finger^b, Raphael Y. de Camargo^c, Fernando Castor Filho^d, Fábio M. Costa^e

^a*Federal University of Maranhão, Department of Informatics, São Luís, Brazil*

^b*University of São Paulo, Department of Computer Science, São Paulo, Brazil*

^c*Federal University of ABC, Santo André, Brazil*

^d*Federal University of Pernambuco, Informatics Center, Recife, Brazil*

^e*Federal University of Goiás, Institute of Informatics, Goiânia, Brazil*

Abstract

The InteGrade project is a multi-university effort to build a novel grid computing middleware based on the opportunistic use of resources belonging to user workstations. The InteGrade middleware currently enables the execution of sequential, bag-of-tasks, and parallel applications that follow the BSP or MPI programming models.

This article presents the lessons learned over the last five years of InteGrade development and describes the achieved solutions concerning the support for robust application execution. The contributions cover the related fields of application scheduling, execution management, and fault tolerance. We present our solutions, describing their implementation principles and evaluation through the analysis of several experimental results.

Key words: grid computing, opportunistic grid, resource management, fault tolerance

Email address: `fssilva@deinf.ufma.br` (Francisco José da Silva e Silva)

Preprint submitted to Journal of Parallel and Distributed Computing February 3, 2010

1. Introduction

The success of grid systems can be verified by the increasing number of middleware systems, actual production grids, and dedicated forums that appeared in recent years. The use of Grid Computing technology is increasing rapidly, reaching more scientific fields and encompassing a growing body of applications (Wilkinson , 2009; Grandinetti , 2005).

A grid might be seen as a way to interconnect clusters that is much more convenient than the construction of huge clusters. Another possible approach for conceiving a grid is the opportunistic use of workstations of regular users. The focus of an opportunistic grid middleware is not on the integration of dedicated computer clusters (e.g., Beowulf) or supercomputing resources, but on taking advantage of idle computing cycles of regular computers and workstations that can be spread across several administrative domains.

During the last five years, our research group has been engaged on the development of the InteGrade project¹, a multi-university effort to build a robust and flexible middleware for opportunistic grid computing. InteGrade's main goal is to be an opportunistic grid environment with support for tightly-coupled parallel applications. By leveraging the idle computing power of existing commodity workstations and connecting them to a grid infrastructure, InteGrade enables the execution of computationally-intensive parallel applications that would otherwise require expensive cluster or parallel machines.

In this article, we focus on the support provided by InteGrade for applica-

¹Homepage: <http://www.integrade.org.br>

tion execution, covering three key related issues concerning the development of an opportunistic grid middleware: resource management and availability prediction for application scheduling, execution management, and fault tolerance.

Resource management encompasses challenges such as how to efficiently monitor a large number of highly distributed computing resources belonging to multiple administrative domains. On opportunistic grids, this issue is even harder due to the dynamic nature of the execution environment, where nodes can join and leave the grid at any time due to the use of the non-dedicated machines by their regular (non-grid) users. An effective monitoring infrastructure is crucial for performing appropriate application scheduling decisions and for timely detecting failures. Besides the current state of grid resources, the application scheduler should also take into consideration a prediction of the future availability of computing resources. This is particularly useful on opportunistic grids, as users of non-dedicated machines can resume local processing, forcing grid tasks to either migrate to other grid machine or abort and possibly restart at another machine.

With respect to application **execution management**, which also includes monitoring, there must be user-friendly mechanisms to execute applications in the grid environment, to control the execution of jobs, and to provide tools to collect application results and to generate reports about current and past situations. Application execution management should encompass all execution models supported by the middleware. The InteGrade middleware currently allows the execution of sequential, bag-of-tasks, and parallel applications that follow the Bulk Synchronous Parallel (BSP) (Valiant , 1990)

or Message Passing Interface (MPI) (MPI Forum , 1997) programming models. A related issue concerning application execution is the management of application data, which includes the application binaries, and input and output data. On grid environments, applications usually generate large amounts of data, which, combined with the environment's large scale, turns a centralized approach for data storage inappropriate. On the other hand, building a flexible distributed storage system that provides high data availability and fast data access in a dynamic environment, such as an opportunistic grid, is a challenging task.

Finally, **fault tolerance** comprises a major requirement for grid middleware as grid environments are highly prone to failures, a characteristic amplified on opportunistic grids due their dynamism and the use of non-dedicated machines, leading to a non-controlled computing environment. An efficient and scalable failure detection mechanism must be provided by the grid middleware, along with a means for automatic application execution recovery, without requiring human intervention.

InteGrade presents solutions to those problems. There are special modules in charge of controlling the resources for opportunistic grids, monitoring them, and providing tools to analyze the gathered data. There are special tools for managing the execution of applications, which guarantee secure mechanisms to store application binaries, their input and the generated output. InteGrade also provides mechanisms for fault tolerance, including both replication and checkpointing, in a configurable way, allowing multiple choices for data storage. The middleware also provides support for tightly-coupled parallel applications using both the MPI and BSP standard

programming models, as well as mobile agents for long-running sequential applications.

All these solutions, when used together, provide a powerful framework that allows for the execution of sequential and parallel applications in an opportunistic environment. Indeed, as all these solutions were designed to be decentralized and distributed, a good level of scalability is obtained.

This article concentrates on the new developments on the InteGrade project carried out since its prototype architecture was presented five years ago; for a detailed description of that early work, the reader should refer to (Goldchleger et al. , 2004).

In the next section we present an overview of the current InteGrade architectural model. In Section 3, we discuss the use of resource monitoring data to infer future availability of resources, which is taken into account during the scheduling process to minimize application migration and restart. Section 4 describes aspects related to application execution management, emphasizing the mechanisms that allow the execution of parallel applications that follow the MPI or BSP models; this section also addresses the management of application input and output data. In Section 5, we describe the available mechanisms used to guarantee application progress, circumventing failures. In Section 6, we provide a comparative analysis of InteGrade and Condor, a well know middleware system for opportunistic grids. Finally, in Section 7 we conclude the article presenting open research problems and ongoing work.

2. InteGrade Overview

The basic architectural unit of an InteGrade grid is a cluster, a collection of machines usually connected by a local network. Clusters can be organized in a hierarchy, enabling the construction of grids with a large number of machines. Each cluster contains a *Cluster Manager* node that hosts InteGrade components responsible for managing cluster resources and for inter-cluster communication. Other cluster nodes are called *Resource Providers* and export part of their resources to the grid. They can be either shared with local users (e.g., secretaries using a word processor) or dedicated machines. The cluster manager node, containing InteGrade management components, must be a stable machine, usually a server, but not necessarily dedicated to InteGrade execution only. In case of a cluster manager failure, only its managed cluster machines will become unavailable.

InteGrade currently allows the execution of three application classes: (a) sequential applications, where the task to be run is assigned to a single grid node²; (b) parametric or bag-of-tasks applications, where several copies of a task are assigned to different grid nodes, each of them processing a subset of the input data independently and without exchanging data; (c) parallel applications following the BSP or MPI models, whose processes exchange data among themselves using message passing or shared memory abstractions.

The InteGrade architecture comprises several components. Figure 1 shows the components that enable application execution, which are the following.

Application Repository (AR): before being executed, an application

²A machine acting as a grid resource provider.

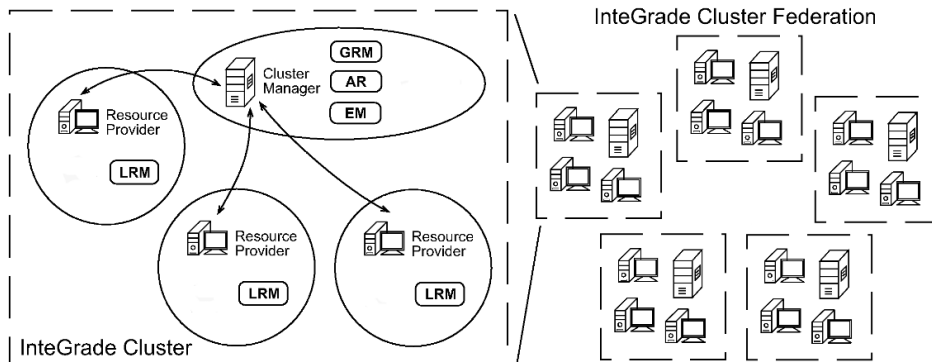


Figure 1: InteGrade architecture.

must be previously registered with the Application Repository. This component stores the application description (metadata) and binary code.

Application Submission and Control Tool (ASCT): a graphical user interface that allows users to browse the content of the Application Repository, submit applications, and control their execution. Alternatively, applications can be submitted via the **InteGrade Grid Portal**, a Web interface similar to ASCT.

Local Resource Manager (LRM): a component that runs on each cluster node, collecting information about the state of resources such as memory, CPU, disk, and network. It is also responsible for instantiating and executing applications scheduled to the node.

Global Resource Manager (GRM): manages cluster resources by receiving notifications of resource usage from the LRMs in the cluster (through an information update protocol) and runs the scheduler that allocates tasks to nodes based on resource availability; it is also responsible for communication with GRMs in other clusters, allowing applications to be scheduled for execution in different clusters. Each cluster has a GRM and, collectively,

the GRMs form the Global Resource Management service. We assume that the cluster manager node where the GRM is instantiated has a valid IP address and firewalls are configured to allow TCP traffic on the port used by the GRM. Network administrators establishing a Virtual Organization can, optionally, make use of ssh tunnels in order to circumvent firewalls and NAT boxes.

Execution Manager (EM): maintains information about each application submission, such as its state, executing node(s), input and output parameters, submission and termination timestamps. It also coordinates the recovery process in case of application failures.

Since grids are inherently more vulnerable to security threats than traditional systems, as they potentially encompass a large number of users, resources, and applications managed by different administrative domains, InteGrade encompass an opinion-based grid security model called Xenia. Xenia provides an authorization and authentication system and a security API that allows developers to access a security infrastructure that provides facilities such as digital signatures, cryptography, resource access control and access rights delegation. Using Xenia, we developed a secure Application Repository infrastructure, which provides authentication, secure communication, authorization, and application validation. A more detailed description of InteGrade security infrastructure can be found on (Pinheiro Junior, Vidal, Kon and Finger , 2006).

From this basic description, we can move on to InteGrade's core issues related to application execution, which include resource management and availability prediction for an appropriate scheduling, application execution

management and fault tolerance, which are explained in the following sections.

3. Resource Availability Prediction

The success of an opportunistic grid depends on a good scheduler. An idle machine is available for grid processing, but whenever its local users need their resources back, grid applications executing at that machine must either migrate to another grid machine or abort and possibly restart at another machine. In both cases, there is considerable loss of efficiency for grid applications. A solution is to avoid such interruptions by scheduling grid tasks on machines that are expected to remain idle for the duration of the task.

InteGrade predicts each machine's idle periods by locally performing Use Pattern Analysis of machine resources at each machine on the grid, as described in (Finger, Bezerra and Conde , 2008, 2009). Currently, four different types resource are monitored: CPU use, RAM availability, disk space, and swap space.

Use pattern analysis deals with *machine resource use objects*. Each object is a vector of values representing the time series of a machine's resource use, as illustrated in Figure 1. The sampling of a machine's resource use is performed at a fixed rate (currently, once every 5 minutes) and grouped into objects covering 48 hours with a 24-hour overlap between consecutive objects. We employ 48-hour long objects so as to have enough past information to be used in the runtime prediction phase.

Use Pattern Analysis performs unsupervised machine learning (Barlow , 1999; Theodoridi and Koutroumba , 2003) to obtain a fixed number of *use*

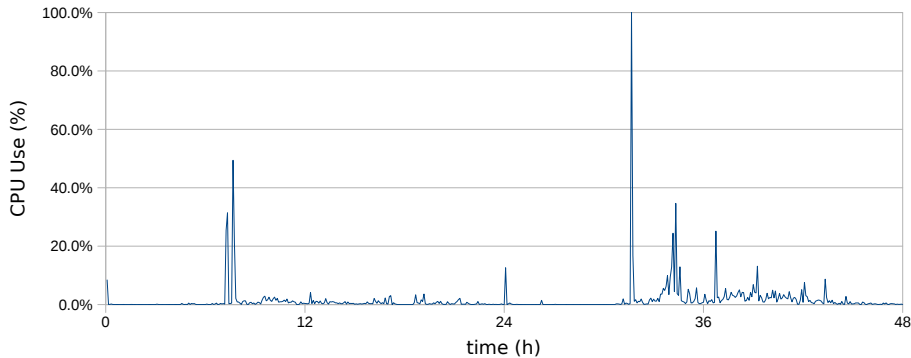


Figure 2: An object representing a machine’s CPU use over a 48h period

classes, where each class is represented by its *prototypical* object. The idea is that each class represents a frequent use pattern, such as a busy work day, a light work day or a holiday. As in most machine learning processes, there are two phases involved in the process, which in the InteGrade architecture are implemented by a module called Local Use Pattern Analyzer (LUPA), as follows.

The Learning Phase. Learning is performed off-line, using 60 objects collected by LUPA during the machine regular use. A clustering algorithm (Sokal , 1996; Everitt, Landau and Leese , 2001) is applied to the training data, such that each cluster corresponds to a use class, represented by a prototypical object, which is obtained by averaging over the elements of the class. Learning can occur only when there is a considerable mass of data. In our case, we require at least two months of data. As data collection proceeds, more data and more representative classes are obtained.

The Decision Phase. There is one LUPA module per machine on the grid. Requests are sent by the scheduler specifying the amount of resources

(CPU, disk space, RAM, etc.) and the expected duration needed by an application to be executed at that machine. The LUPA module decides whether this machine will be available for the expected duration, as explained below. LUPA is constantly keeping track of the current use of resources. For each resource, it focuses on the recent history, usually the last 24 hours, as illustrated in Figure 3, and computes a distance between the recent history and each of the use classes learned during the training phase. This distance takes into account the time of the day in which the request was made, so that the recent history is compared to the corresponding times in the use classes. The class with the smallest distance is the *current use class*, which is used to predict the availability in the near future. If all resources are predicted to be available, then the application is scheduled to be executed; otherwise, it is rejected.

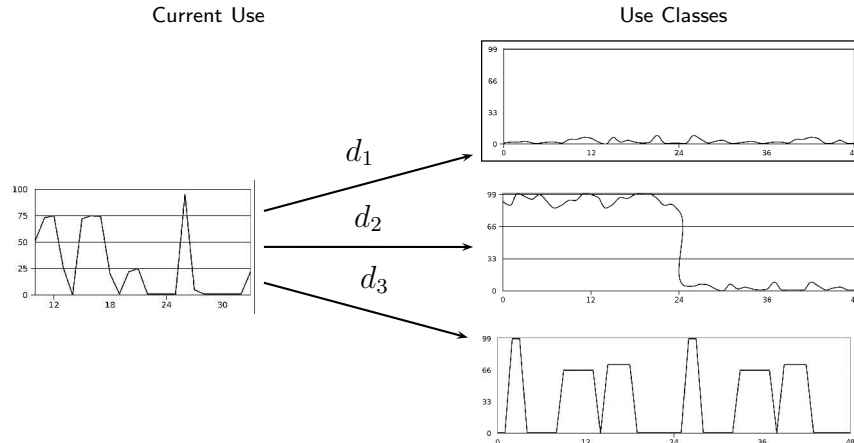


Figure 3: Estimation of class pertinence for use prediction

This scheduling strategy has been tested via simulation, using data collected from real machines on the grid, and was compared against two other

strategies, namely: one that projects the availability at the time of the request into the future (`alt1`); and one that projects into the future a fixed value obtained from the average use of the past 24 hours (`alt2`).

Table 1 presents a summary of the trace-based simulation results with data collected over a period of 7 months. Results are presented for two kinds of machines, single-user and multi-user machines, which have quite different use profiles; and for two types of resource availability prediction: memory and CPU. At random instants of time, a request is issued for a random amount of resources during a random amount of time, and a prediction for availability is made at that time; the prediction of availability is considered correct if the amount of resource required remains available for the full duration of the request; a prediction of unavailability is considered correct if at some point of this interval the requested amount of resources is unavailable. Three basic results are shown. First, the percentage of correct predictions. Second, the probability that the proposed strategy predicts availability with higher percentage of correctness than strategy `alt1`; this probability is computed assuming that the correctness of predictions follows a normal distribution for each strategy, such that the difference between them also follows a normal distribution. Third, the probability that the proposed strategy performs better than strategy `alt2` under the same assumptions. All results in Table 1 were obtained with 5 data clusters. The same experiments were run for 10 clusters, with basically the same results.

Table 1 allows us to conclude that the proposed strategy is very successful in predicting idleness, with success rates above 90% at all times. It also shows that single-user machines are easier to predict than multi-user machines. The

Resource	Machine	Correct Predictions	Beats alt1	Beats alt2
Memory	single user	99.47%	47.77%	99.12%
Memory	multi user	99.08%	64.97%	90.37%
CPU	single user	97.57%	50.21%	87.89%
CPU	multi user	92.52%	91.23%	91.35%

Table 1: Simulation results for 5 clusters

proposed method also performs better than the proposed alternatives in all cases except one, namely the prediction of memory availability for single user machines, in which case the extremely simple strategy of projecting into the future the current availability seems to perform better.

A complete description of Use Pattern Analysis method can be found in our previous work (Finger, Bezerra and Conde , 2008). Work in progress is addressing the problem of predicting, for each application, its resource demands and duration. Admittedly, it is not reasonable to assume that users or even application writers can always perform well such prediction task. So automated methods are under study. Those methods involve computing several statistics over runs of the same program with a range of input data. Several (mainly conservative) prediction methods are being considered, such as the worst case seen in practice, or the average plus three times the standard deviation.

4. Managing Application Executions

Once scheduled to proper grid resources, the tasks that comprise a submitted application can be instantiated and begin their execution. Executing

computationally intensive parallel applications on dynamic heterogeneous environments, such as computational grids, is a daunting task. This is particularly true when using non-dedicated resources, as in the case of opportunistic computing, where one uses only the idle periods of the shared machines. In this scenario, the execution environment is typically highly dynamic, with resources periodically leaving and joining the grid. When a resource becomes unavailable, due to a failure or simply because the machine owner requests its use, the system needs to perform the necessary steps to restart the tasks on different machines. In the case of BSP or MPI parallel applications, the problem is even worse, since all processes that comprise the application may need to be restarted from a consistent distributed checkpoint.

The Execution Manager (EM) module is in charge of managing application execution, including sequential, bag-of-tasks, MPI, and BSP applications. The EM maintains information regarding each application executing on its cluster, such as the execution status, the nodes where the application processes are running and the names of the application input and output files. The EM is also responsible for restarting applications that had processes executing on a machine that became unavailable. Tasks comprising a BSP or MPI parallel application are always scheduled to a single InteGrade cluster in order to minimize communication overhead and avoid the necessity of bypassing firewalls and NAT boxes.

The following subsections describe how InteGrade implements the support for managing the execution of MPI and BSP parallel applications. Next, we describe the management of application data, which includes the application binaries, input, and output data.

4.1. Managing BSP Applications

InteGrade's support for executing BSP applications adheres to the Oxford BSP API ³, targeted for the C language. Thus, an application based on the Oxford BSPLib can be executed over InteGrade with little or even no modification of its source code, requiring only its recompilation and linkage with the appropriate InteGrade libraries.

A BSP computation proceeds in a series of global supersteps. Each superstep comprises three ordered stages: (1) *concurrent computation*: computations take place on every participating process. Each process only uses values stored on its local memory. Computations are independent in the sense that they occur asynchronously of all others; (2) *communication*: at this stage, the processes exchange data between themselves; (3) *barrier synchronization*: when a process reaches this point (the barrier), it waits until all other processes have finished their communication actions. The synchronization barrier is the end of a superstep and the beginning of another one.

InteGrade's implementation of the BSP model uses CORBA (OMG , 2008) for inter-process communication. CORBA has the advantage of being an easier and cleaner communication environment, shortening development and maintenance time and facilitating system evolution. Also, since it is based on a binary protocol, the performance of CORBA-based communication is an order of magnitude faster than the performance of technologies based on XML, requiring less network bandwidth and processing power. On the shared grid machines, InteGrade uses OiL ((Maia, Cerqueira and Cosme

³The Oxford BSP Toolset <http://www.bsp-worldwide.org/implmnts/oxtool>

, 2006)), a very light-weight version of a CORBA ORB that imposes a small memory footprint. Nevertheless, CORBA usage is completely transparent to the InteGrade application developer, who only uses the BSP interface (Goldchleger et al. , 2005).

InteGrade's BSPLib associates to each process of a parallel application a *BspProxy*. The *BspProxy* is a CORBA servant responsible for receiving related communications from other processes, such as a virtual shared address read or write, or the receipt of messages signaling the end of the synchronization barrier. The creation of *BspProxies* is entirely handled by the library and is totally transparent to users.

The first created process of a parallel application is called *Process Zero*. Process Zero is responsible for assigning an unique identifier to each application process, broadcasting the CORBA IORs of each process to allow them to communicate directly, and coordinating synchronization barriers. Moreover, *Process Zero* executes its normal computation on behalf of the parallel application.

On InteGrade, the synchronization barriers of the BSP model are used to store checkpoints during execution, since they provide global, consistent points for application recovery. In this way, in the case of failures, it is possible to recover application execution from a previous checkpoint, which can be stored in a distributed way as described in Section 5.2. Application recovery is also available for sequential, bag-of-tasks, and MPI applications.

4.2. Managing MPI Applications

Support for parallel applications based on MPI is achieved in InteGrade through MPICH-IG (Cardozo and Costa , 2008), which in turn is based on

MPICH2⁴, an open source implementation of the second version of the MPI standard, MPI2 (MPI Forum , 1997). MPICH-IG adapts MPICH2 to use InteGrade's LRM and EM instead of the MPI daemon (MPD) to launch and manage MPI applications. It also uses the application repository to retrieve the binaries of MPI applications, which are dynamically deployed just prior to launch, instead of requiring them to be deployed in advance, as with MPICH2. MPI applications can thus be dispatched and managed in the same way as BSP or sequential applications.

In order to adapt MPICH2 to run on InteGrade, two of its interfaces were re-implemented: the *Channel Interface* (CI) and the *Process Management Interface* (PMI). The former is required to monitor a sockets channel to detect and treat failures. The latter is necessary to couple the management of MPI applications with InteGrade's Execution Manager (EM), adding functions for process location and synchronization.

Regarding communication among an application's tasks, MPICH-IG uses the MPICH2 Abstract Device Interface (ADI) to abstract away the details of the actual communications mechanisms, enabling higher layers of the communications infrastructure to be independent from them. In this way, we implemented two underlying communications channels: a CORBA-based one, for tasks running on different, possibly heterogeneous, networks, and a more efficient one, based on sockets, for tasks that reside in a single cluster.

Another feature of MPICH-IG, in contrast with conventional MPI platforms, refers to the recovery of individual application tasks after failures.

⁴<http://www.mcs.anl.gov/research/projects/mpich2/>

This prevents the whole application from being restarted from scratch, thus contributing to reduce the makespan of application execution. Application recovery is supported by monitoring the execution state of tasks, as described in Section 5.1. Faulty tasks are then resumed on different grid nodes, selected by the GRM. Task recovery is implemented using checkpoints, which are managed as described in Section 5.2, although using system-level checkpoints, which are more suitable for tightly coupled parallel applications. While other MPI platforms focus specifically on fault-tolerance and recovery, notably MPICH-V (Bosilca et al. , 2002), they usually rely on homogeneous, dedicated, clusters. MPICH-IG removes this limitation to enable the dynamic scheduling of non-dedicated machines.

These features also favor MPICH-IG when compared to other approaches to integrate MPI into grid computing environments, such as MPICH-G2 (Karonis, Toonen and Foster , 2003). In common to MPICH-G2 is the ability to run MPI applications on large scale heterogeneous environments, as well as the ability to switch from one communications protocol to another, depending on the relative location of the application tasks. However, MPICH-IG's ability to use non-dedicated resources in an opportunistic way further contributes to scale up the amount of available resources. In addition, MPICH-IG enables legacy MPI applications to be transparently deployed on an InteGrade grid, without the need to modify their source code.

Regarding performance, Figure 4 shows a comparison of MPICH-IG (without checkpointing) and MPICH2, considering the execution time for a parallel QuickSort algorithm. The execution time is shown for different input sizes (number of elements to be sorted). As can be seen MPICH-IG has compara-

ble performance for small input sizes. Actually, its performance in such cases was slightly worse due to the overhead of loading application binaries from the application repository prior to execution on each node (in MPI, application binaries are assumed to be pre-installed on the nodes). However, as the input size grows, it shows a considerable gain. This is due to InteGrade’s ability to schedule tasks to the most available nodes, while MPICH2’s MPD simply chooses the next node from a circular list.

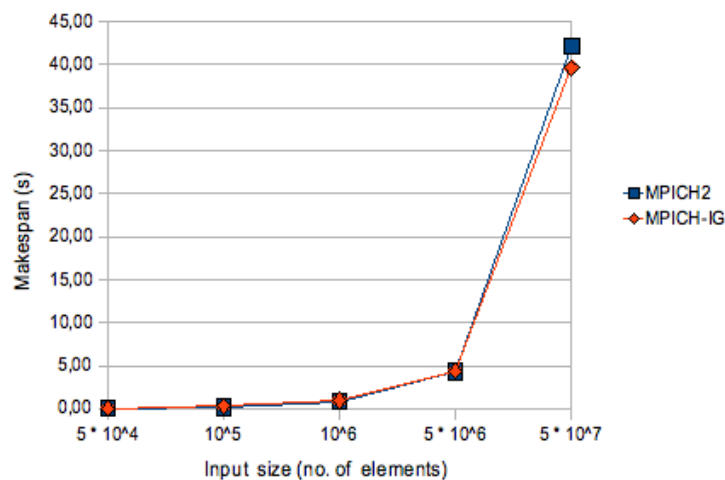


Figure 4: Comparing the execution time for parallel QuickSort on MPICH-IG and MPICH2, with 10 parallel tasks and 6 processing nodes.

4.3. Managing Application Data

Users can register an application in InteGrade using either the Application Submission and Control Tool (ASCT) or the InteGrade Web portal, providing the application description and one or more execution files to enable its execution on multiple platforms. InteGrade stores the application executables in the Application Repository (AR) module, together with meta-

data describing the application and the platforms for which executables are available.

Managing and storing application input and output files is more difficult than managing executables, since the former can be much larger. To deal with those files, we developed a distributed data repository called OppStore (Camargo and Kon , 2007). Access to this distributed repository is performed through a library called *access broker*, which interacts with OppStore.

OppStore is a middleware that provides reliable distributed data storage using free disk space from shared grid machines. The goal is to use this free disk space in an opportunistic way, i.e., only during the idle periods of the machines. The system is structured as a federation of clusters and is connected by a Pastry peer-to-peer network (Rowstron, Druschel , 2001) in a scalable and fault-tolerant way. This federation structure allows the system to disperse application data throughout the grid. During storage, the system slices the data into several redundant, encoded fragments and stores them in different grid clusters. This distribution improves data availability and fault-tolerance, since fragments are located in geographically dispersed clusters. When performing data retrieval, applications can simultaneously download file fragments stored in the highest bandwidth clusters, enabling efficient data retrieval.

When an InteGrade user submits an application for execution, its input files are first stored in distributed grid machines using the access broker. The execution request is then sent to the GRM, which selects the LRMs that will execute the application and forwards the request to those nodes. When a

LRM receives an application execution request, it obtains the application binary for its particular platform from the Application Repository and the input file from OppStore. When the execution finishes, the output files are stored in OppStore.

Using OppStore, application input and output files can be obtained from any node in the system. Consequently, after a failure, restarting of an application execution in another machine can be easily performed. Also, when an application execution finishes, the output files uploaded to the distributed repositories can be accessed by the user from any machine connected to the grid.

To determine the availability of data stored in a large-scale grid composed of non-dedicated machines, we simulated a grid composed of 100 clusters, with the number of machines on each cluster randomly chosen as 10, 20, 50, 100, and 200. We defined three usage patterns, based on measurements of machine utilization in different environments (Mutka and Livny , 1991; Bolosky, Douceur, Ely and Theimer , 2000), and which are randomly assigned to each cluster. In the first pattern, the mean idle time is 60% during the day and 80% during the night and weekends. The second pattern has idle times of 25% and 40%, and the third 40% and 70%, respectively. We distributed the clusters uniformly across 24 time zones.

We simulated the storage of ten thousand files and attempted to later retrieve those files, checking the number of fragments that could be recovered for each file. The file retrieval requests were repeated during a simulated period of one month. Figure 5 shows the percentage of successful retrieval requests for the file retrieval attempts. Using only the idle periods of the

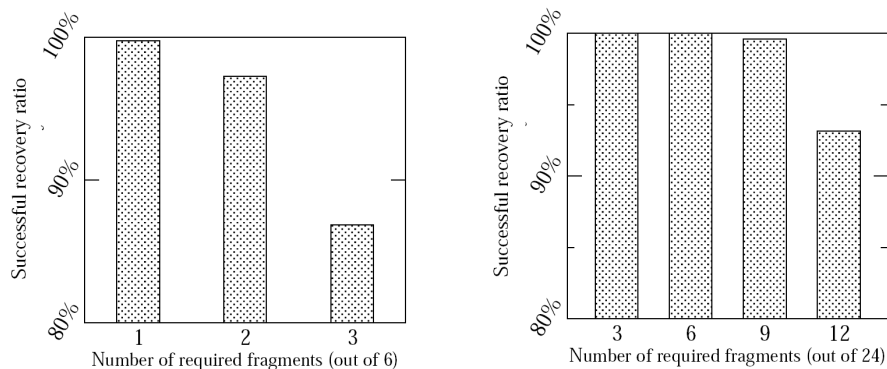


Figure 5: Percentage of successful file retrieval requests.

shared machines to retrieve data in realistic situations, we could recover enough fragments to reconstruct the files in 99.9% of the requests, for files encoded in 24 fragments, from which 8 are required for reconstruction, a replication factor of 3. When all the fragments necessary to reconstruct a file are not available immediately, the broker can download the available fragments and wait until the remaining fragments become available.

We also performed experimental evaluations (Camargo and Kon , 2007) that showed that the time required to store and retrieve files in OppStore is mostly dependent on the bandwidth of cluster connections. The broker downloads the file fragments in parallel and, since the number of available fragments is normally larger than the number of required fragments for file reconstruction, the broker can choose the fragments located in the clusters with the fastest connections. Consequently, file retrieval can be performed very efficiently in OppStore.

5. Fault Tolerance in Grid Environments

On opportunistic grids, application execution can fail due to several reasons. System failures can result not only from an error on a single component but also from the usually complex interactions between the several grid components that comprise a range of different services. In addition to that, grid environments are extremely dynamic, with components joining and leaving the system at all times. Also, the likelihood of errors occurring during the execution of an application is exacerbated by the fact that many grid applications will perform long tasks that may require several days of computation.

To provide the necessary fault tolerance functionality for grid environments, several services must be available, such as: (a) **failure detection**: grid nodes and applications must be constantly monitored by a failure detection service; (b) **application failure handling**: various failure handling strategies can be employed in grid environments to ensure the continuity of application execution; and (c) **stable storage**: execution states that allow recovering the pre-failure state of applications must be saved in a data repository that can survive grid node failures. This section describes InteGrade failure detection and handling techniques. InteGrade stable storage is based on Oppstore, already presented on Section 4.3.

5.1. *InteGrade Failure Detection*

Failure detection is a very important service for large-scale opportunistic grids. The very high rate of churn makes failures a frequent event and the capability of the grid infrastructure to efficiently deal with them has a direct impact on its ability to make progress. Hence, failed nodes should be

detected quickly and the monitoring network should itself be reliable, so as to ensure that a node failure does not go undetected. At the same time, due to the scale and geographic dispersion of grid nodes, failure detectors should be capable of disseminating information about failed nodes as fast and reliably as possible and work correctly even when no process has a globally consistent view of the system. Moreover, the non-dedicated nature of opportunistic grids requires that solutions for failure detection be very lightweight in terms of network bandwidth consumption and usage of memory and CPU cycles of resource provider machines. Besides all of these requirements pertaining to the functioning of failure detectors, they must also be easy to set-up and use; otherwise they might be a source of design and configuration errors. It is well-known that configuration errors are a common cause of Grid failures (Medeiros et al. , 2003).

The aforementioned requirements are hard to meet as a whole and, to the best of our knowledge, no existing work in the literature addresses all of them. This is not surprising, as some goals, e.g., a reliable monitoring network and low network bandwidth consumption, are inherently conflicting. Nevertheless, they are all real issues that appear in large-scale opportunistic grids, and reliable grid applications are expected to deal with them in a realistic setting. The InteGrade failure detection service (Filho et al. , 2008) is an attempt to achieve this goal. It includes a number of features that, when combined and appropriately tuned, address all these challenges while adopting reasonable compromises for the ones that conflict. The most noteworthy features of the proposed failure detector are the following: (i) a gossip- or infection-style approach (Renesse, Minsky and Hayden , 1998),

meaning that the network load imposed by the failure detector scales well with the number of processes in the network and that the monitoring network is highly reliable and decentralized; (ii) self-adaptation and self-organization in the face of changing network conditions; (iii) a crash-recover failure model, instead of simple crash; (iv) ease of use and configuration; (v) low resource consumption (memory, CPU cycles, and network bandwidth).

InteGrade's failure detection service is completely decentralized and runs on every grid node. Each process in the monitoring network established by the failure detection service is monitored by K other processes, where K is an administrator-defined parameter. This means that for a process failure to go undetected, all the K processes monitoring it would need to fail at the same time. Even for high process failure probabilities, e.g., 0.2, a small value of K , such as 6 or 7, makes this event unlikely. A process r which is monitored by a process s has an open TCP connection with it through which it sends heartbeat messages and other kinds of information. If r perceives that it is being monitored by more than K processes, it cancels the monitoring relationship with a randomly chosen process, so as to keep the number of monitoring connections as close as possible to K . Furthermore, r also periodically checks whether the number of processes monitoring is lower than K . If so, r asks a randomly chosen process to monitor it. It does this periodically, so long as the sum of the number of processes monitoring r and the number of processes that r has asked to monitor it (but have still not answered) is lower than K . Evaluation of this simple approach has shown that it yields considerable gains in reliability (Filho et al. , 2009) at a very low cost in terms of extra control messages.

InteGrade’s failure detector automatically adapts to changing network conditions. Instead of using a fixed timeout to determine the failure of a process, it continuously outputs the probability that a process has failed based on the inter-arrival times of the last W heartbeats and the time elapsed since the last heartbeat was received, where W is an administrator-defined parameter. The failure detector can then be configured to take recovery actions whenever the failure probability reaches a certain threshold. Multiple thresholds can be set, each one triggering a different recovery action, depending on the application requirements.

We employ a reactive and explicit approach to disseminate information about failed processes. This means that once a process learns about a new failure it automatically sends this information to J randomly-chosen processes that it monitors or that monitor it. The administrator-defined parameter J dictates the speed of dissemination. According to (Ganesh et al. , 2003), for a system with N processes, if each process disseminates a piece of information to $(\log n) + c$ randomly chosen processes, the probability that the information does not reach every process in the system is $e^{-e^{(-c)}}$, with $n \rightarrow \infty$. For $J = 7$, this probability is less than 0.001. On the other hand, no explicit action is taken to disseminate information about new processes. Instead, processes get to know about new processes by simply receiving heartbeat messages. Each heartbeat that a process p sends to a process q includes some randomly chosen ids of K processes that p knows about. In a grid, it is important to quickly disseminate information about failed processes in order to initiate recovery as soon as possible and, when recovery is not possible and the application has to be re-initiated, to avoid wasting grid resources.

On the other hand, information about new members is not so urgent, since not knowing about new members in general does not keep grid applications from making process.

Our group membership service is implemented in Lua (Ierusalimschy , 1996), an extensible and lightweight programming language. Lua makes it easy to use the proposed service from programs written in other programming languages, such as Java, C, and C++. Moreover, it executes in several platforms. Currently, we have successfully run the failure detector in Windows XP, Mac OS X, and several flavors of Linux. The entire implementation of the group membership service comprises approximately 80Kb of Lua source code, including comments.

5.2. *InteGrade Application Recovery*

In order to overcome application execution failures, InteGrade provides support for the most used failure handling strategies: (1) **retrying**: when an application execution fails, it is restarted from scratch; (2) **replication**: the same application is submitted for execution multiple times, generating various application replicas; all replicas are active and execute the same code with the same input parameters at different nodes; and (3) **checkpointing**: periodically saves the process' state in stable storage during the failure-free execution time. Upon a failure, the process restarts from the latest available saved checkpoint, thereby reducing the amount of lost computation. As part of the application submission process, users can select the desired technique to be applied in case of failure. These techniques can also be combined resulting in four more elaborate failure handling techniques: *retrying* (without checkpoint and replication), *checkpointing* (without replication), *replication*

(without checkpointing), and *replication with checkpointing*.

The InteGrade **checkpointing** mechanism adds to the basic application execution protocol steps responsible for gathering application state into a checkpoint and its storage in machines executing the Autonomous Data Repository (ADRs) module of OppStore. ADRs are located at resource provider machines in the grid and are responsible for managing grid data stored on these machines.

To gather application state, InteGrade includes a portable application-level checkpointing mechanism (Camargo, Kon and Goldman , 2005) for sequential, bag-of-tasks, and BSP parallel applications written in C. This portability allows an application's stored state to be recovered on a machine with a different architecture from the one where the checkpoint was generated. In application-level checkpointing, the application is responsible for providing the data that will be checkpointed. We implemented a pre-compiler that inserts, into application code, the statements responsible for gathering and restoring the application state from the checkpoint. To use checkpointing, the user has to instrument the application source code using the precompiler and then recompile the application using his preferred compiler. On BSP applications, checkpoints are generated immediately after the end of a BSP synchronization phase. Since all messages are guaranteed to be delivered during that phase, there is no risk of message losses or processes receiving extra messages. For MPI parallel applications, we provide a system-level checkpointing mechanism based on a coordinated protocol (Cardozo and Costa , 2008).

When the checkpointing library needs to store a checkpoint, it request

a list of available ADRs from OppStore. When storing checkpoints, OppStore selects only ADRs from the cluster where the application is running. Comparing with the standard storage mode of OppStore, called perennial mode and described in Section 4.3, checkpoint storage has lower overhead, since all machines are physically close, but reduced fault-tolerance, since it is vulnerable to cluster failures. A trade-off is achieved by storing part of the checkpoints in the perennial mode. When storing checkpoints in the local cluster, OppStore allows two storage strategies:

- *Data replication*: stores full copies of checkpoint data on different machines. This approach uses more storage space and network bandwidth, but less processing power, since no coding is performed;
- *Rabin's classic Information Dispersal Algorithm (IDA)* (Rabin , 1989): encodes the checkpoint into redundant fragments, such that regenerating the original checkpoint is possible using only a subset of them. This approach uses less network bandwidth and storage space, but more processing power.

These two strategies allow the application library to select a trade-off between the use of storage space and network bandwidth and the CPU time necessary to encode the checkpoint. Checkpoint encoding and transfer are performed by a separate application thread, allowing the application to concurrently continue its execution. Finally, to improve performance, the library transfers the several checkpoint fragments or replicas to the repositories in parallel.

We evaluated the overhead of the checkpointing mechanism using the

replication and IDA strategies using a matrix multiplication (Hayashida et al. , 2005) application. Figure 6 shows the results for the executions of the application using two matrix sizes (3200x3200 and 4800x4800), configured to run without checkpointing and with checkpointing using replication and two IDA configurations. The first, IDA(7,1), generates 8 fragments from which 7 are required to reconstruct the original file, and the second, IDA(6,2), generates 8 fragments and require 6 for file reconstruction. We executed each experiment 16 times, with the vertical bars showing the mean measured overhead and the error bars representing the standard deviation. For a matrix of size 4800x4800, which generates global checkpoints of 791MB, and a minimum interval between checkpoints of 60s, the overhead is 13% when using replication, 18% when using IDA(7,1) and 20% for IDA(6,2). This overhead can be reduced by increasing the checkpointing interval. For example, with a checkpointing interval of 5 minutes, the overhead would be of approximately 2.6% using replication.

These results show that the use of checkpointing has low overhead even when running coupled parallel applications that generate large checkpoints. Using the IDA strategy causes a higher overhead, as it is necessary to encode the checkpoints using processor cycles from the machines running the parallel application. Using replication causes a smaller overhead, but uses more network and storage resources. Since parallel applications run on grid clusters composed of physically close machines and storage space is not a scarce resource, the default strategy used is replication of the checkpoint data.

Replication of the execution of application tasks is another failure handling technique commonly applied on grid environments. While checkpoint-

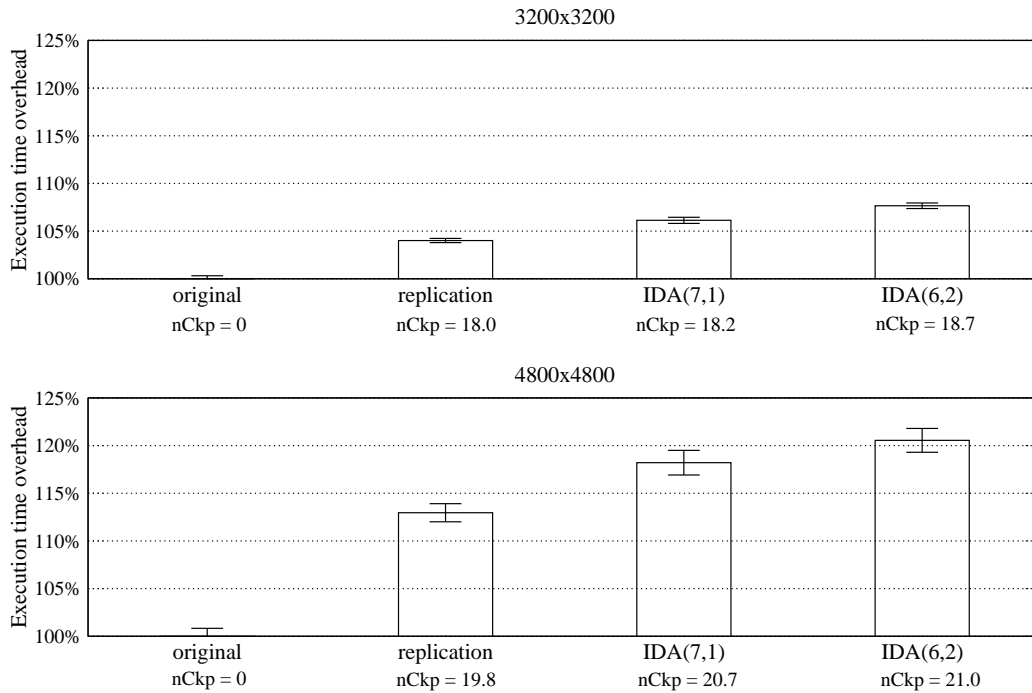


Figure 6: Checkpointing execution overhead using different checkpointing strategies.

ing imposes an overhead to application execution (for gathering and storing the execution state), replication requires a larger amount of grid resources for executing the same task.

InteGrade allows replication for sequential, bag-of-tasks, MPI and BSP applications. The amount of generated replicas is currently defined during the application execution request, issued through the Application Submission and Control Tool (ASCT). The request is forwarded to the Global Resource Manager (GRM), which runs a scheduling algorithm that guarantees that all replicas will be assigned to different nodes. Another InteGrade component, called Application Replication Manager (ARM), concentrates most of the code responsible for managing replication. The GRM instantiates a new

ARM for every application execution request that demands replication. The ARM gathers the application input files from the ASCT and forwards a copy of the execution request to the LRM of every node involved with the application execution. Each LRM instantiates an application process. The ARM also registers each application process with the Execution Manager (EM) and, in case of a replica failure, starts its recovery process. When the first application replica concludes its job, the ARM kills the remaining ones, releasing the allocated grid resources. In this way, the complexity of replication is confined to the ARM code, minimizing the required changes on InteGrade components to support replication.

6. Condor and InteGrade: a Comparative Analysis

A well known middleware system that allows the construction of opportunistic grids is Condor (Litzkow, Livny and Mutka , 1988; Thain, Tannenbaum and Livny , 2005). The project started in 1988 and currently has a mature implementation and a large number of users. It has some similarities to InteGrade, such as the support for the opportunistic execution of application in shared machines. It supports the execution of sequential and parallel applications, including Parallel Virtual Machine (PVM) and MPI applications. Scheduling in Condor is performed using a mechanism called ClassAds, which tries to match advertised resources and application needs. It also has a checkpointing mechanism for execution fault-tolerance, but it works for sequential applications only.

The main difference between InteGrade and Condor is that InteGrade focuses on the opportunistic usage of resources for the execution of coupled

parallel applications. InteGrade has two mechanisms for allowing this execution: (1) a usage pattern analyzer that predicts the amount of time that each machine is likely to remain idle, and (2) a checkpoint-based rollback recovery mechanism to allow failed BSP and MPI applications to be restarted from the last saved checkpoint. Both mechanisms are central in the InteGrade design and are a major difference between the systems.

7. Summary, Conclusions, and Future Directions

Opportunistic grid middleware enables the use of the existing computing infrastructure available in laboratories and offices in universities, research institutes, and companies to execute computationally intensive parallel applications. Nevertheless, executing this class of applications on such a dynamic and heterogeneous environment is a daunting task, especially when non-dedicated resources are used, as in the case of opportunistic computing.

This article presented recent advances on the InteGrade opportunistic grid middleware concerning the support for application execution, covering the related fields of application scheduling, execution management, and fault tolerance.

The InteGrade module called Local Usage Pattern Analyzer (LUPA) predicts each machine's idle periods by locally performing Usage Pattern Analysis of machine resources using an unsupervised machine learning approach. Experimental results demonstrate that the proposed strategy is very successful in predicting idleness, with success rates above 90% at all times, leading to better scheduling decisions and minimizing task migrations.

InteGrade is very flexible with respect to the supported application classes.

It currently allows the execution of sequential, parametric (bag-of-tasks), and parallel applications following either the BSP or MPI models. Native MPI applications can be transparently deployed on an InteGrade grid, without the need to modify their source code or even to recompile them. InteGrade's MPI support emphasizes scalability as well as the opportunistic use of idle resources, which are not normally feasible with cluster-based MPI platforms. It also allows the recovery of individual tasks of an application, which prevents the whole application from being restarted from scratch. InteGrade also allows the execution of BSP applications through the implementation of several functions from the Oxford BSPLib, including both DRMA (direct remote memory access) and BSMP (bulk synchronous message passing). Fault-tolerance for BSP applications is provided by an application-level checkpointing mechanisms, which includes a precompiler that instruments the application source code.

Concerning the management of application data, which includes the application binaries, input and output data, InteGrade's OppStore component provides reliable distributed data storage using the free disk space from shared grid machines in an opportunistic way, i.e., only during the idle periods of the machines. Experimental results demonstrate that OppStore can recover enough fragments to reconstruct files in 99.9% of the tested requests and that file retrieval can be performed very efficiently.

Since opportunistic grid environments are highly prone to failures, special care was taken on InteGrade to circumvent application execution disruptions. The failure detection infrastructure separates failure detection from the dissemination of information about failures, allowing the grid adminis-

trator to adjust two distinct parameters: the number of failure detectors monitoring a node and the number of notified failure detectors. A gossip-or infection-style approach is also used, leading to a scalable, reliable, and decentralized monitoring approach. InteGrade provides support for replication, checkpointing, and retry recovery techniques that can also be combined together in a flexible way. Checkpointing and replication are available for sequential, bag-of-tasks, MPI and BSP parallel applications and experimental results demonstrate that checkpoint usage has a low overhead even when running coupled applications that generate large checkpoints.

All those presented features, when combined, provide a robust and flexible environment that allows the efficient execution of large scale, computationally-intensive parallel applications, even on highly dynamic environments, such as opportunistic grids. As future work we intend to deploy the InteGrade middleware in a large opportunistic grid and evaluate its performance in a real scenario with hundreds of machines and users. The InteGrade source code and documentation is available as free software distributed under the LGPL license at <http://www.integrate.org.br>.

References

- H. B. Barlow. Unsupervised learning. In G. Hinton and T. J. Sejnowski, editors, *Unsupervised Learning: Foundations of Neural Computation*, pages 1–17. MIT Press, 1999.
- W. J. Bolosky, J. R. Douceur, D. Ely, and M. Theimer. Feasibility of a serverless distributed file system deployed on an existing set of desktop PCs. *SIGMETRICS Performance Evaluation Review*, 28(1):34–43, 2000.

- G. Bosilca, A. Bouteiller, F. Cappello, S. Djilali, G. Fedak, C. Germain, T. Herault, P. Lemarinier, O. Lodygensky, F. Magniette, V. Neri, and A. Selikhov. MPICH-V: toward a scalable fault tolerant MPI for volatile nodes In Supercomputing '02: Proceedings of the 2002 ACM/IEEE Conference on Supercomputing, pages 1–18, Baltimore, Maryland, 2002. IEEE Computer Society Press.
- R. Y. de Camargo and F. Kon. Design and implementation of a middleware for data storage in opportunistic grids. In *CCGrid '07: Proceedings of the 7th IEEE/ACM International Symposium on Cluster Computing and the Grid*, Washington, DC, USA, 2007. IEEE Computer Society.
- R. Y. de Camargo, F. Kon, and A. Goldman. Portable checkpointing and communication for BSP applications on dynamic heterogeneous grid environments. In *SBAC-PAD'05: The 17th International Symposium on Computer Architecture and High Performance Computing*, Rio de Janeiro, Brazil, October 2005.
- M. C. Cardozo, F. M. Costa. MPI support on opportunistic grids based on the InteGrade middleware In *Proceedings of the 2nd Latin American Grid International Workshop (LAGrid)*, Campo Grande, Brazil, 2008.
- B. Everitt, S. Landau, and M. Leese. *Cluster Analysis*. A Hodder Arnold Publication, 4th edition edition, 2001.
- F. C. Filho, R. Castro, A. Marques, F. Soares-Neto, R. Y. de Camargo, and F. Kon. A robust and scalable group membership service for large-scale

- grids (*in portuguese*). In *Proceedings of the SBRC'2009 Workshop on Grid Computing and Applications*, Recife, Brazil, 2009.
- F. C. Filho, A. Marques, R. Y. de Camargo and F. Kon. A group membership service for large-scale grids. In *Proceedings of the 6th International Workshop on Middleware for Grid Computing*, Leuven, Belgium, 2008.
- M. Finger, G. C. Bezerra, and D. M. R. Conde. Resource use pattern analysis for opportunistic grids. In *6th International Workshop on Middleware for Grid Computing (MGC 2008)*, Leuven, Belgium, December, 2008.
- M. Finger, G. C. Bezerra, and D. M. R. Conde. Resource use pattern analysis for predicting resource availability in opportunistic grids. *Concurrency and Computation: Practice and Experience*, Accepted, 2009.
- A. J. Ganesh, A.-M. Kermarrec, and L. Massoulié. Peer-to-peer membership management for gossip-based protocols. *IEEE Transactions on Computers*, 52(2):139–149, Feb. 2003.
- A. Goldchleger, A. Goldman, U. Hayashida, F. Kon. The implementation of the BSP parallel computing model on the InteGrade grid middleware In *Proceedings of the 3rd International Workshop on Middleware for Grid Computing*, Grenoble, France, 2005. ACM Press.
- A. Goldchleger, F. Kon, A. Goldman, M. Finger, and G. C. Bezerra. InteGrade: object-oriented grid middleware leveraging idle computing power of desktop machines. *Concurrency and Computation: Practice and Experience*, 16(5):449–459, March 2004.

- L. Grandinetti. *Grid Computing: The New Frontier of High Performance Computing*. Elsevier, 2005.
- U. K. Hayashida, K. Okuda, J. Panetta and S. W. Song. Generating parallel algorithms for cluster and grid computing. In *Proceedings of the the 2005 International Conference on Computational Science and its Applications*, pages 509–516, May 2005.
- R. Ierusalimschy, L. H. de Figueiredo, and W. C. Filho. Lua - an extensible extension language. *Software: Practice Experience*, 26(6):635–652, 1996.
- J. de R. B. P. Júnior, A. C. T. Vidal, F. Kon, and M. Finger. Trust in large-scale computational grids: An SPKI/SDSI extension for representing opinion. In *4th International Workshop on Middleware for Grid Computing - MGC 2006*, Melbourne, Australia, November 2006.
- N. Karonis, B. Toonen, and I. Foster. MPICH-G2: a grid-enabled implementation of the message passing interface. *Journal of Parallel and Distributed Computing (JPDC)*, 63(3):551–563, May 2003.
- M. Litzkow, M. Livny, and M. Mutka. Condor - A hunter of idle workstations. *ICDCS '88: Proceedings of the 8th International Conference of Distributed Computing Systems*, 104–111, June 1988.
- R. Maia, R. Cerqueira, and R. Cosme. OiL: an object request broker in the Lua. In *Proc. 5th Tools Session of the Brazilian Symposium on Computer Networks (SBRC2006)*, Curitiba, Brazil, June 2006.
- R. Medeiros, W. Cirne, F. V. Brasileiro, and J. P. Sauv e. Faults in grids: why are they so bad and what can be done about it? In *Proceedings of*

- the 4th IEEE International Workshop on Grid Computing (GRID 2003)*, pages 18–24, Phoenix, USA, November 2003.
- MPI Forum. MPI-2: extensions to the Message-Passing Interface. MPI Forum. <http://www.mpi-forum.org/docs/mpi-20-html/mpi2-report.html> (accessed on 02/17/2009), July 1997.
- M. W. Mutka and M. Livny. The available capacity of a privately owned workstation environment. *Performance Evaluation*, 12(4):269–284, 1991.
- Object Management Group (OMG). Common Object Request Broker Architecture (CORBA) Specification, Version 3.1. Available at <http://www.omg.org/spec/CORBA/3.1/>. 2008.
- M. O. Rabin. Efficient dispersal of information for security, load balancing, and fault tolerance. *Journal of ACM*, 36(2):335–348, 1989.
- R. van Renesse, Y. Minsky, and M. Hayden. A gossip-style failure detection service. In *Proceedings of Middleware'1998*, Lake District, England, September 1998.
- A. I. T. Rowstron and P. Druschel. Pastry: scalable, decentralized object location, and routing for large-scale peer-to-peer systems In *Proceedings of the Middleware 2001: IFIP/ACM International Conference on Distributed Systems Platforms*, pages 329–350, isbn 3-540-42800-3, Heidelberg, Germany, 2001
- R. R. Sokal. Clustering and classification: background and current directions. In *In Proceedings of the Advanced Seminar on Classification and Clustering*, 1996.

- D. Thain, T. Tannenbaum, and M. Livny. Distributed computing in practice: the Condor experience: research articles. *Concurrency and Computation: Practice and Experience*, 17(2-4):323–356, 2005.
- S. Theodoridis and K. Koutroumba. *Pattern Recognition*. Academic Press, 2003.
- L. G. Valiant. A bridging model for parallel computation. *Communications of the ACM*, 33(8):103–111, 1990.
- B. Wilkinson. *Grid Computing: Techniques and Applications*. Chapman & Hall/CRC, 2009.