

SPECIAL ISSUE PAPER

GRIDHPC: A decentralized environment for high performance computing

Bilal Fakh¹  | Didier El Baz^{1,2} | Igor Kotenko²

¹LAAS-CNRS, Université de Toulouse,
Toulouse, France

²ITMO University, Saint Petersburg, Russia

Correspondence

Bilal Fakh, LAAS-CNRS, Université de
Toulouse, 31400 Toulouse, France.
Email: bfakh@laas.fr

Funding information

Agence Nationale de la Recherche;
Government of Russian Federation,
Grant/Award Number: 08-08

Summary

This paper presents GRIDHPC, a decentralized environment dedicated to high performance computing. It relies on the reconfigurable multi network protocol RMNP to support data exchange between computing nodes on multi network systems with Ethernet, Infiniband, Myrinet, and on OpenMP for the exploitation of computing resources of multicore CPU. We report on scalability of several parallel iterative schemes of computation combined with GRIDHPC. In particular, the experimental results show that GRIDHPC scales up when combined with asynchronous iterative schemes of computation.

KEYWORDS

asynchronous iterations, computing environment, grid computing, heterogeneous networks, high performance computing, loosely synchronous applications

1 | INTRODUCTION

In this paper, we present the decentralized environment GRIDHPC dedicated to High Performance Computing (HPC) on grid platforms. The HPC applications that we consider are basically loosely synchronous applications like the solution of numerical simulation problems¹ that present frequent data exchanges between computing nodes. GRIDHPC facilitates the use of large scale distributed systems and the work of programmer. In particular, it uses a limited number of communication operations.

The GRIDHPC environment allows data exchange between computing nodes with multi network multi-core configurations. It relies on the Reconfigurable Multi Network Protocol (RMNP) to support data exchange on multi-network systems and on OpenMP² for the exploitation of computing resources of multi-core CPU.

The protocol of communication RMNP is an extension of the Configurable Transport Protocol (CTP)³ that makes use of the Cactus framework.⁴ RMNP can configure itself automatically and dynamically in function of application requirements like scheme of computation that is implemented, ie, synchronous or asynchronous iterative schemes and elements of context like available network interface cards and network topology by choosing the most appropriate communication network and mode between computing nodes. It can use simultaneously several networks like Ethernet, Infiniband, and Myrinet. These features are particularly important since we consider loosely synchronous applications that present frequent data exchanges between computing nodes. To the best of our knowledge, these features have not been carried out previously on environments or runtime systems in the literature.

The remainder of this paper is organized as follows. Related work is presented in Section 2. Section 3 deals with our contribution to support communication in a multi-network context. In particular, the Reconfigurable Multi Network Protocol RMNP is presented. Section 4 presents the architecture and task assignation of the GRIDHPC environment. Parallel programming model is given in Section 5. Computational results with the decentralized environment GRIDHPC for the obstacle problem are displayed and analyzed in Section 6. Section 7 concludes this paper.

2 | RELATED WORK

The possibility to consider heterogeneous network resources for HPC applications goes back to Madeleine and MPICH-Madeleine.^{5,6} Recently, research has focused on runtime systems for HPC applications carried out on heterogeneous architectures that combine multi-core CPU and computing accelerators.

Heterogeneous multi-core platforms mixing CPUs and computing accelerators are nowadays widely spread. High Performance ParalleX (HPX)^{7,8} is a C++ runtime system for parallel and distributed applications, some of which are loosely synchronous applications. It has been developed for systems of any scale and aims to address issues like resiliency and power efficiency. It has a programming model unifying all types of parallelism available in HPC systems which uses the available resources to attain scalability. It is portable and easy to use. It is published under an open-source license and has an active user community. It is built using dynamic and static data flow, fine grain future-based synchronization, and continuation style programming. The main goal of HPX is to create an open source implementation of the ParalleX execution model⁹ for conventional systems like classic Linux based Beowulf clusters, Android, Windows, Macintosh, Xeon/Phi, Blue Gene/Q, or multi-socket highly parallel SMP nodes.

HPXCL¹⁰ and APEX¹¹ are libraries, which provide additional functionality that extend the HPX. HPXCL allows programmers to incorporate GPUs into their HPX applications. Users write an OpenCL kernel and pass it to HPXCL, which manages the synchronization and data offloading of the results with the parallel execution flow on the CPUs. APEX gathers arbitrary information about the system and uses it to make runtime-adaptive decisions based on user-defined policies.

StarPU¹² is a runtime system that provides an interface to execute parallel tasks over heterogeneous hardware (multi-core processors and computing accelerators) and develop easily powerful scheduling algorithms. It is based on the integration of the data-management facility with a task execution engine. The main components of StarPU are a software distributed shared memory (DSM) and a scheduling framework; DSM enables task computations to overlap and avoid redundant memory transfers. The scheduling framework maintains an up-to-date and a self-tuned database of kernel performance models over the available computing tasks to guide the task mapping algorithms. We note that middle layers tools like programming environments and HPC libraries can build up on top of StarPU to allow programmers to make existing applications exploit different accelerators with limited effort.

To the best of our knowledge, the aforementioned runtime systems do not handle multi-network contexts that we find typically in grids. In the next sections, we show how we have addressed this issue and we propose a decentralized environment for multi-network and multi-core grid platforms.

3 | COMMUNICATION PROTOCOL

This section presents the mechanisms that support multi-network communications.

3.1 | Heterogeneous multi-cluster environment

Figure 1 displays a multi-network platform with interconnected stand-alone computing nodes that can work cooperatively as a single integrated computing resource. In particular, Figure 1 shows the architecture of typical sets of computing nodes built around low-latency high bandwidth

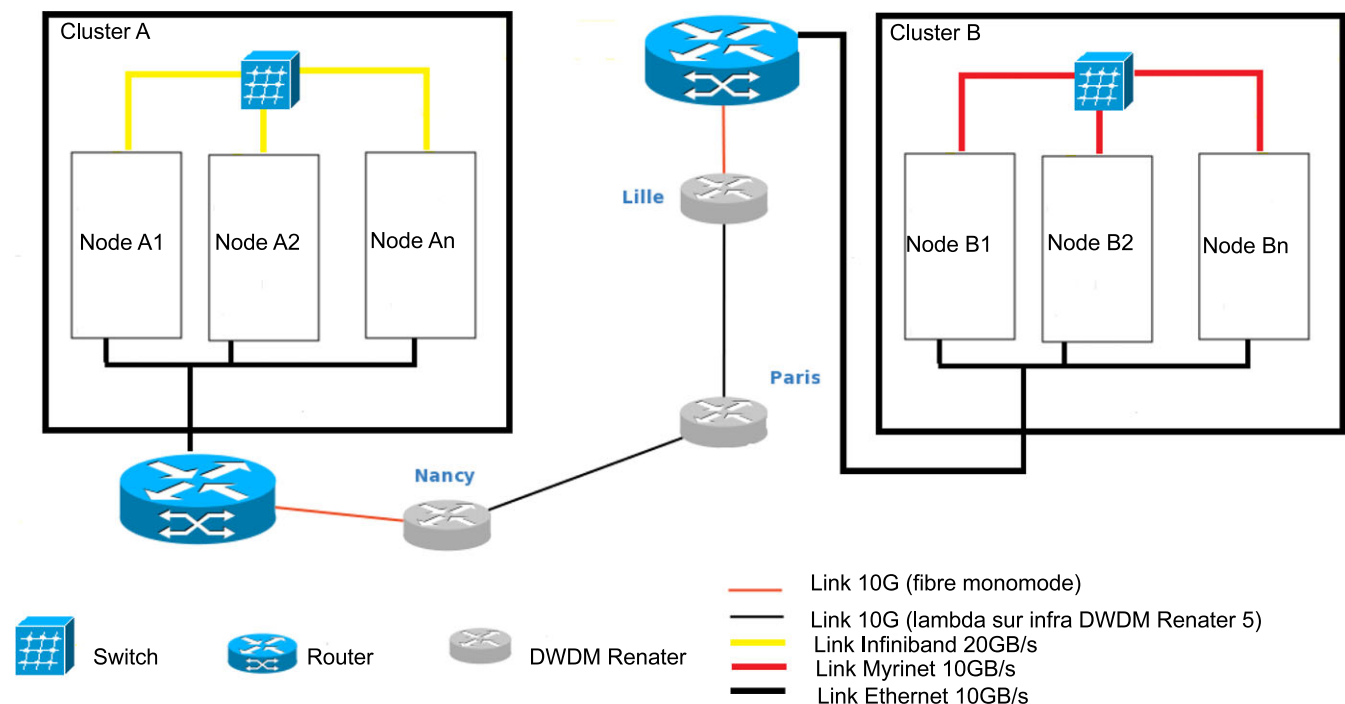



FIGURE 1 Example of multi-network distributed platform

TABLE 1 Example of content of Htable and test on the location of the hosts thanks to the comparison of IP addresses

	IP addresses of computing node A1	IP addresses of computing node A2	IP addresses of computing node B1	IP addresses of computing node B2
Nic Ethernet	192.16.64.10	192.16.64.11	192.16.34.20	192.16.34.21
Nic Infiniband Or Nic Myrinet	192.18.64.10	192.18.64.11	192.18.34.20	192.18.34.21



 Test on the third group of the IP addresses of two hosts
 in order to determine if they belong to the same cluster

interconnection network like Infiniband and Myrinet. Supporting heterogeneous networks mainly consists in integrating a functionality in order to switch from one network to another, according to the communication needs.

3.2 | Htable

3.2.1 | Distance metric

We consider a distance metric that is based on IP address. In particular, it concentrates on the third group of the IP address. For example, in the case of three computing nodes, ie, A1 having IP address 192.16.64.10, A2 having IP address 192.16.64.11, and B1 having IP address 192.16.34.20, the value of third group of A1 and A2 is equal to 64, while the value of third group of B1 is 34. Consequently, it is deduced that A1 and A2 have the same location (they are in the same cluster), while A1 and B1 are in different locations (see Table 1).

3.2.2 | Definition of Htable

Htable is designed to manage several network adapters within the same application session. In particular, Htable permits each computing node to switch between the networks according to the communication needs.

Several network interface cards (NICs) are added to the interface of the RMNP communication protocol and information about these NICs are stored in the Htable. In the Htable (see Table 1), the IP addresses displayed in the first line correspond to Ethernet network and the IP addresses given in the second line, if any, correspond to fast network like Infiniband or Myrinet.

3.3 | Reconfigurable multi-network protocol RMNP

The Reconfigurable Multi-Network Protocol RMNP aims at enabling an efficient use of the complete set of underlying communication software and hardware available in a given multi-network system that uses, eg, Ethernet, Infiniband, and Myrinet. It is able to deal with several networks via the management of several networks adapters.

3.3.1 | Choice of networks

The network management procedure has two steps (see Algorithm 1). The first step corresponds to the test of the locality between the computing nodes thanks to the comparison of IP addresses and the second step corresponds to the choice of the appropriate network for data exchange depending on the locality of computing nodes. In particular, the second step is based on choosing the best interface network (high bandwidth and low latency network) from the Htable according to the result of the locality test. Consequently, if the locality test returns that the considered computing nodes have different locations, then the Ethernet network interface is chosen to perform the communication between the two computing nodes. If the locality test returns that the computing nodes have the same location, then the best network interface in the Htable is selected, eg, Infiniband or Myrinet.

3.3.2 | Example of scenario

We present now a simple scenario for the RMNP communication protocol so as to illustrate its behavior. We consider a high performance computing application, like for instance a large scale numerical simulation application, solved on the network composed of two clusters shown in Figure 1. Computing nodes in cluster A own both a Fast-Ethernet card, ie, 192.16.64.x, and Infiniband card, ie, 192.18.64.x (see Table 1) and computing nodes in cluster B own both a Fast-Ethernet card, ie, 192.16.34.x, and Myrinet card, ie, 192.18.34.x where x is a value between 1 and 255. We suppose that we have a communication network between computing nodes like $A1 \leftrightarrow A2$, $A2 \leftrightarrow B1$, and $B1 \leftrightarrow B2$, where $X \leftrightarrow Y$

means that there is bidirectional link between X and Y. The value of third group of A1 and A2 is equal to 64. Consequently, the communication between A1 and A2 that share the same high speed networks is made via Infiniband network. The values of third group of A2 and B1 are 64 and 34, respectively. Hence, the communication between the considered computing nodes is made via Ethernet network. The value of third group of B1 and B2 is equal to 34. Consequently, the communication between the considered computing nodes is made via Myrinet network.

Algorithm 1 Test locality and choose the best network

function Localityandbestnet (laddress, raddress);

Input : laddress and raddress : Ethernet IP addresses of local and remote computing nodes

/* inet_ntop converts the network address structure into a character string */

inet_ntop(AF_INET, &(laddress.sin_addr), l, 80);

inet_ntop(AF_INET, &(raddress.sin_addr), r, 80);

substringl ← GetValThirdGroup(l);

/* GetValThirdGroup function get the value of third group of a given IP address */

substringr ← GetValThirdGroup(r);

if strcmp(substringl, substringr) == 0 **then**

/* Local and Remote computing nodes belong to the same cluster */

Bestladdress ← Get(laddress);

/* Get function get the best interface network from Htable */

Bestraddress ← Get(raddress);

/* Communication between computing nodes are made via the best network (Infiniband or Myrinet Network), ie, Bestladdress and Bestraddress */

else

/* Local and Remote computing nodes are in different cluster, Hence communication between them are made via Ethernet Network, ie, laddress and raddress */

4 | THE DECENTRALIZED ENVIRONMENT GRIDHPC

4.1 | Environment architecture of GRIDHPC

The decentralized environment GRIDHPC natively supports any combination of networks and multi-core CPUs by using the reconfigurable multi network protocol RMNP and OpenMP. Figure 2 shows the architecture of GRIDHPC. It consists of five main components.¹³

4.1.1 | Interface environment component

It is the interaction interface between the application like obstacle problem and the environment. It allows users to submit their tasks and retrieve final results.

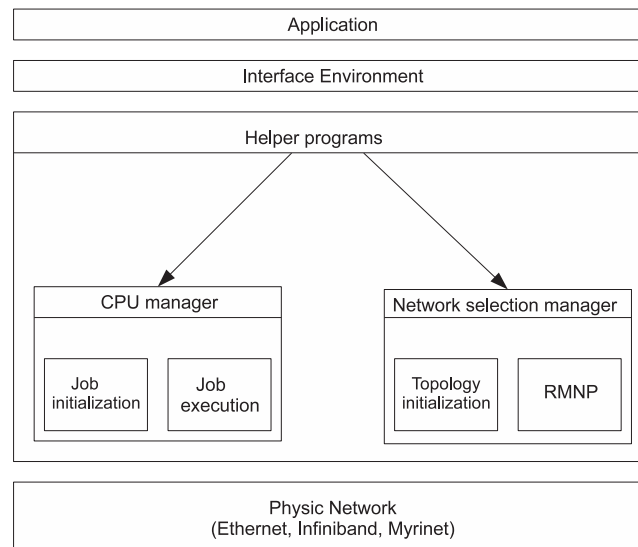


FIGURE 2 Environment architecture of GRIDHPC

4.1.2 | Helper programs

GRIDHPC works with tools called helper programs that are responsible for the analysis of the application, task assignation, and building the network topology. The helper programs rely on two pillars, namely, the CPU manager and the Network selection manager (see Figure 2). The CPU manager is composed of Job Initialization and Job Execution. The Network selection manager is composed of Topology Initialization and RMNP-OpenMP components. In the CPU manager module, there are two components.

- Job Initialization Component is responsible for problem decomposition and assignment of tasks to individual CPU cores.
- Job Execution Component executes sub-tasks on the different CPU cores and takes care of data exchange, ie, communication of updates produced by the parallel iterative method. At the end of the application, it regroups the results from all the computing cores.

In the Network selection manager module, there are two components.

- Topology Initialization Component organizes connected computing nodes into clusters and maintains links between clusters. It is based on storing in the Htable information regarding the network interface cards (NIC) used in the application by the different computing nodes.
- RMNP Component provides support for directed data exchange between computing nodes on several networks like Infiniband, Myrinet, and fast Ethernet using the reconfigurable multi-network protocol RMNP.

We note that the CPU manager is in charge of data exchange between computing cores, ie, read/write, while the network selection manager is in charge of data exchange between computing nodes via the best underlying network, ie, high speed and low latency network like Infiniband and Myrinet. The combination of the CPU manager and the Network selection manager permits us to use the decentralized environment GRIDHPC in a multi-network and multi-core context.

4.2 | Processor hierarchy and GRIDHPC

Task assignation in GRIDHPC¹³ is based on a hierarchical Master-Worker paradigm that relies on three entities, ie, a master, several sub-masters (coordinators), and several workers.

The master or submitter is the unique entry point; it gets the entire application as a single original task, ie, root task. The master decomposes the root task into sub-tasks and distributes these sub-tasks amongst a farm of workers. The master takes also care of gathering the scattered results in order to produce the final result of the computation.

The sub-masters or coordinators are intermediary entities that enhance scalability. They forward sub-tasks from the submitter to workers and return results to the submitter limiting network congestion.

The workers run in a very simple way. They receive a message from the sub-master that contains their assigned sub-tasks and they distribute them to their computing cores. They perform computations and data exchange with neighboring computing nodes. At the end of the application, when the iterative schemes have converged, they regroup the results from all their computing cores and send them back to the coordinator. The coordinators transfer results to the submitter. Note that the number of workers in a group cannot exceed 32 in order to ensure efficient management of a sub-master. Figure 3 shows an example of processor hierarchy.

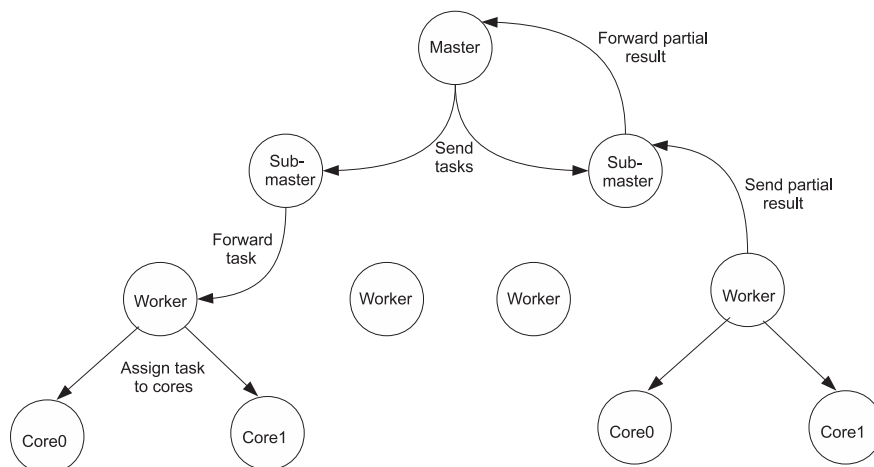


FIGURE 3 Processor hierarchy

5 | PARALLEL PROGRAMMING MODEL OF GRIDHPC

5.1 | Communication operations

We aim at facilitating the use of GRID platforms as well as the programming of large scale HPC applications and hiding the complexity of communication management. The communication protocol RMNP has a reduced set of communication operations; there are only GRID_Send, GRID_Receive, and GRID_Wait. Contrarily to MPI communication library where communication mode is fixed by the semantics of communication operations, the communication mode of a given communication operation depends on the context at application level like chosen parallel iterative scheme of computation, eg, synchronous or asynchronous iterative scheme and elements of context like topology at network level, ie, inter or intracluster communication and type of network like Ethernet, Infiniband, and Myrinet. The programming model permits us to expect scalable performance and application flexibility. The prototype of the communication operations of our programming model are summarized in Listing 1 where:

- **GRID_Send** communication operation is used to send a message placed in buffer to subtask destination.
- **GRID_Receive** communication operation is used to receive a message from sub-task source.
- **GRID_wait** operation is used to wait for a message from another computing node.

Note that flags parameters in these operations are used to distinguish two types of messages, ie, CTRL_FLAG indicates control messages and DATA_FLAG indicates data messages. Data messages are used to exchange updates between computing nodes, while control messages are used to exchange information related to computation state like state of termination condition, termination command, etc. These data are particularly important for the convergence detection process and termination phase.

5.2 | Application programming model

Figure 4 displays the activity diagram that a parallel application must follow. The diagram consists of 13 activities.

- **Task definition:** First, the application is defined at the submitter, ie, setting task parameters as well as computational schemes (synchronous iterations, asynchronous iterations, and hybrid), problem size, and the number of computing nodes required. Note that hybrid iterative scheme is a combination of synchronous and asynchronous computation schemes, ie, synchronous iterations in the same cluster and asynchronous iterations between clusters (at global level).
- **Collect computing nodes:** Based on the task definition, the submitter collects free computing nodes.
- **Enough computing nodes:** The submitter verifies if there are enough free computing nodes to carry out the task. If there are not enough free computing nodes, then the computation is terminated.
- **Send sub-tasks:** If there are enough free computing nodes, then the submitter sends sub-tasks to coordinators.
- **Forward sub-tasks:** The coordinator forwards sub-tasks from submitter to workers.
- **Receive sub-tasks:** Computing nodes receive sub-tasks from coordinator and become workers.
- **Distribute sub-tasks on the different cores:** This decomposes sub-task into sub-sub-tasks and assigns each one to a core at a given computing node. Note that the number of sub-sub-tasks is equal to the maximum number of cores in a computing node.
- **Calculate:** This is the module that performs computations relative to sub-tasks. Each core executes its sub-sub-task. We note that, in the case of applications solved by iterative algorithms, a worker has to carry out many iterations; after each iteration, it has to exchange updates with others workers. For this purpose, it uses RMNP for data exchanges between computing nodes, ie, Grid_Send and Grid_Receive communication operations.
- **Results aggregations of all the cores:** Sub-sub-tasks results are aggregated into one result at a given computing node.
- **Send results:** Sends aggregated results to coordinator.

```

1) int GRID_Send(GRIDSubtask *pSubtask,
    uint32_t dest, char *buffer, size_t
    size, int flags);

2) int GRID_Receive(GRIDSubtask
    *pSubtask, uint32_t source, char
    *buffer, size_t size, int flags);

3) int GRID_Wait(GRIDSubtask* pSubtask,
    uint32_t *iSubtaskRank, int *flags);

```

Listing 1 Prototype of RMNP communication operations

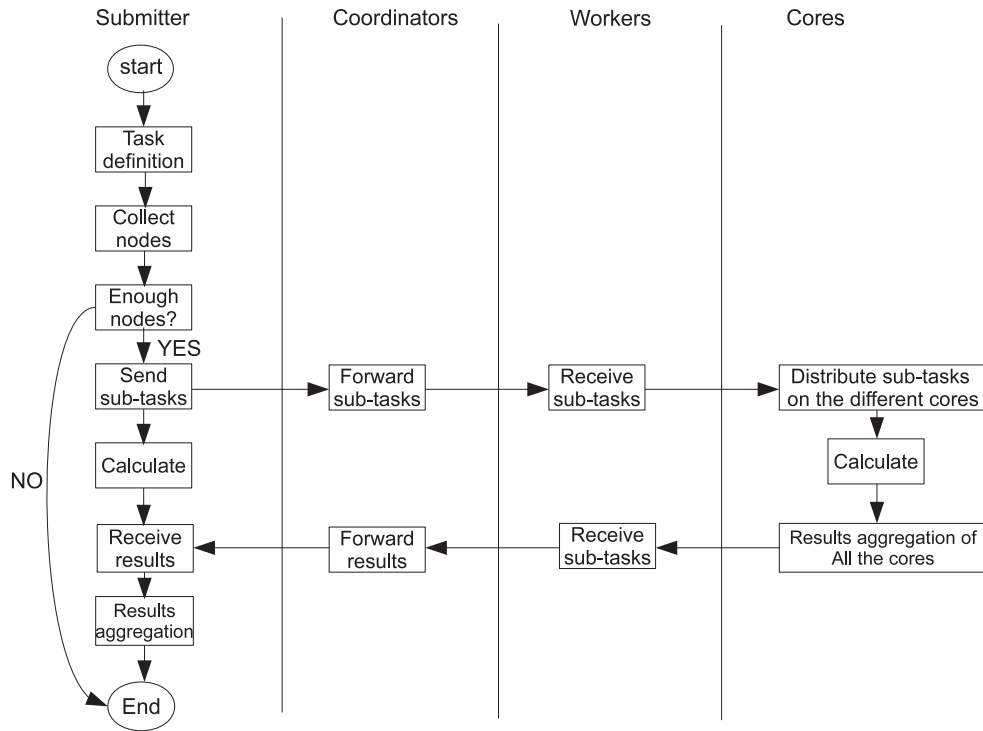


FIGURE 4 Activity diagram of a parallel application with GRIDHPC

- **Forward results:** The coordinator forwards results from workers to submitter.
- **Receive results:** The submitter receives sub-tasks results from coordinators.
- **Results aggregation:** Sub-tasks results of all the workers are aggregated at submitter into final result.

6 | COMPUTING RESULTS AND EVALUATION

This section presents an evaluation of the scalability of GRIDHPC in a multi-core and multi-network context for a loosely synchronous application, ie, the obstacle problem.

6.1 | Obstacle problem

The obstacle problem occurs in many domains like mechanics and finance and can be formulated as follows:

$$\begin{cases} \text{Find } u^* \text{ such that} \\ A.u^* - f \geq 0, u^* \geq \varnothing \text{ everywhere in } \Omega, \\ (B.u^* - f)(\varnothing - u^*) = 0 \text{ everywhere in } \Omega, \\ B.C., \end{cases} \quad (1)$$

where the domain $\Omega \in \mathbb{R}^3$ is an open set, A is an elliptic operator, \varnothing is a given function, and B.C. denotes the boundary conditions.

We consider the discretization of the obstacle problem. The parallel solution of the associated fixed point problem via the projected Richardson method combined with several iterative schemes of computation is studied. Reference is made to the works of Nguyen et al^{14,15} for the mathematical formulation of parallel synchronous and asynchronous projected Richardson methods. The interest of asynchronous iterations for various problems including boundary value problems has been shown in other works.¹⁶⁻¹⁸

The experiments are carried out via GRIDHPC in order to solve the 3D obstacle problem with different parallel iterative schemes of computation, ie, synchronous, asynchronous, and hybrid schemes of computation. We consider cubic domains with $n = 256$, up to 512 points, where n denotes the number of points on each edge of the cube.




FIGURE 6 Termination detection of synchronous iterations

The behavior of computing nodes is given by the finite state machine in Figure 7. It can be summarized as follows: each computing node can have three possible states, ie, Active (A), Inactive (I), and Terminated (T). Four types of messages can be exchanged by computing nodes, ie, activate message, inactivate message, termination message, and update message. Note that the first three message types are control messages, ie, flags = CTRL_FLAG, and the last message type is data message, ie, flags = DATA_FLAG.

In active state (A), a computing node P_r evaluates the local termination test, ie, the local conjunction of all the token of computing cores at P_r ; if it is satisfied, and then P_r does not compute any update; otherwise, each computing core at P_r updates components of the sub-sub-vector assigned to it. After that, P_r aggregates the values of the components of the iterate vector at the interface that were updated by its computing cores and sends them to adjacent computing nodes.

In inactive state (I), a computing node is waiting for messages using GRID_wait operation. Note that, if P_r is inactive and receives an update message from a computing node P_r , then P_r becomes active (A) and it is the children of P_r . If P_r receives an inactivate message from P_r , then P_r removes P_r from its list of children.

Terminated state (T) corresponds to the case where the computation is terminated at the computing node.

To illustrate the procedure, we consider the simple example of the evolution of the activity graph in the case of four computing nodes presented in Figure 8. Initially, only the root, ie, computing node P_1 is active and all other computing nodes are inactive. The computing nodes become progressively active on receiving an update message from other computing nodes. An activity graph is generated; the topology of the graph changes progressively as the messages are received and the local termination tests are satisfied; an active computing node becomes inactive if its list of children is empty and its local termination test is satisfied; then, the computing node sends an inactivate message to its parent. The activity graph changes as the computation progresses. At the end, the computing nodes becomes progressively inactive (the computing node P_1 is the last node to become inactive) and the global termination is detected.

6.3 | Experimental results

This subsection presents an evaluation of GRIDHPC in various multi-core and multi-network contexts for the obstacle problem and several parallel iterative methods. Table 2 gives the characteristics of the different clusters of the Grid5000 platform²⁰ used in the computational experiments.

We study first the scalability of parallel synchronous and asynchronous iterative schemes of computation combined with GRIDHPC for a 3D obstacle problem with size 256^3 . Table 3 displays computing time of parallel iterative schemes of computation for several grid configurations. The synchronous, asynchronous, and hybrid iterative methods are denoted by Syn, Asyn, and Hybrid, respectively. We consider the Graphene cluster, the Chinqchint cluster, and a multi-cluster configuration, ie, Graphene and Chinqchint clusters coupled via 10 Gb/s Ethernet network (measured latency is about 5 microseconds). We note that the environment GRIDHPC selects always the best network in each cluster when several networks are available (Infiniband network with Graphene cluster and Myrinet network with Chinqchint). As an example, synchronous iterations takes 742 seconds with GRIDHPC on Graphene cluster with 128 computing cores, while the same iterative scheme takes 1282 seconds with the same number of computing cores of the same cluster when using Ethernet network.

Table 4 shows the corresponding computing gains. The computing gain is given as follows:

$$\text{computing gain } C_g = t_1/t_s, \quad (2)$$

where t_1 is the fastest parallel computing time on one multi-core machine and t_s is the parallel computing time on several multi-core machines.

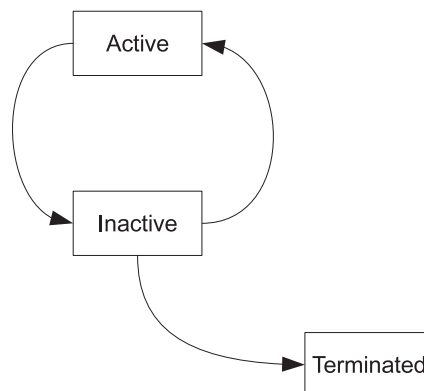



FIGURE 7 States of computing nodes in the termination detection procedure of asynchronous iterations

**FIGURE 8** Evolution of the activity graph**TABLE 2** Characteristics of machines

Site	Cluster	Processors Type	Cores	Interconnection Networks	clock GHz	RAM GB
Lille	Chinqchint	Intel Xeon E5440 QC	8	Ethernet and Myrinet	2.83	8
Nancy	Graphene	Intel Xeon X3440	4	Ethernet and Infiniband	2.53	16
Rennes	Paravance	Intel Xeon E5-2630v3	16	Ethernet 10 Gbs	2.4	128
Grenoble	Edel	Intel Xeon E5520	8	Ethernet and Infiniband	2.27	24
Grenoble	Genepi	Intel Xeon E5420 QC	8	Ethernet and Infiniband	2.5	8

TABLE 3 Computing time of parallel iterative methods applied to the obstacle problem with size 256^3 and several grid configurations

Number of cores	Computing time / s						
	Graphene cluster		Chinqchint cluster		Graphene and Chinqchint clusters		
	Syn	Asyn	Syn	Asyn	Syn	Asyn	Hybrid
4	3115	-	-	-	-	-	-
8	1723	1595	3298	-	-	-	-
16	1355	971	1785	1576	2321	1369	1461
32	1012	594	1419	1001	1824	855	1119
64	841	378	1119	631	1444	536	977
128	742	272	-	-	1262	386	854

TABLE 4 Computing gain of parallel iterative methods applied to the obstacle problem with size 256^3 and several grid configurations

Number of cores	Computing gain						
	Graphene cluster		Chinqchint cluster		Graphene and Chinqchint clusters		
	Syn	Asyn	Syn	Asyn	Syn	Asyn	Hybrid
4	1	-	-	-	-	-	-
8	1.80	1.95	1	-	-	-	-
16	2.29	3.2	1.84	2.09	1.34	2.27	2.13
32	3.07	5.24	2.32	3.29	1.70	3.64	2.78
64	3.7	8.24	2.94	5.22	2.15	5.81	3.18
128	4.19	11.45	-	-	2.46	8.06	3.64

From Tables 3 and 4, we see that, in the case of a single cluster like Graphene or Chinqchint, asynchronous iterative schemes of computation perform better than synchronous iterative schemes since there are no idle time due to synchronization or synchronization overhead. We note also that parallel asynchronous iterations combined with GRIDHPC scale up. The multi-core multi-network configuration considered in Tables 3 and 4 corresponds to the case where Chinqchint cluster in Lille is connected to Graphene cluster in Nancy. Lille and Nancy are two French cities three hundred kilometers apart. The experiments are carried out with up to 24 computing nodes (16 computing nodes at Graphene and 8 computing nodes at Chinqchint) and 128 cores. There is the same number of computing cores in the two clusters, ie, 64 cores at Graphene and 64 cores at Chinqchint. Computing results show that, even in a heterogeneous context where the computing nodes have different number of cores and there are several networks, the combination of GRIDHPC and asynchronous or hybrid iterative schemes of computation scales up. The computing gain of hybrid iterations is situated in between the computing gains of synchronous and asynchronous iterations. This is due to the fact that hybrid schemes of computation correspond to synchronous iterations in the same cluster and asynchronous iterations between clusters.

TABLE 5 Computing time of parallel iterative methods applied to the obstacle problem with size 256^3 on a grid with Edel and Genepi clusters

Number of cores	Computing time / s		
	Edel and Genepi clusters		
	Syn	Asyn	Hybrid
8	1962	-	-
16	1867	1448	1494
32	1571	909	1283
64	1252	538	1132
128	1117	343	1049
256	1020	307	967

TABLE 6 Computing gain of parallel iterative methods applied to the obstacle problem with size 256^3 on a grid with Edel and Genepi clusters

Number of cores	Computing gain		
	Edel and Genepi clusters		
	Syn	Asyn	Hybrid
8	1	-	-
16	1.05	1.35	1.31
32	1.24	2.15	1.52
64	1.56	3.64	1.73
128	1.75	5.72	1.87
256	1.92	6.39	2.02

TABLE 7 Computing time of parallel iterative methods applied to the obstacle problem with size 384^3 , 448^3 and 512^3 on Paravance cluster

Number of cores	Computing time / s					
	384		448		512	
	Syn	Asyn	Syn	Asyn	Syn	Asyn
16	4200	-	8943	-	21705	-
32	4124	2242	9353	5190	21787	13344
64	3317	1785	6690	3507	11798	6643
128	2912	1099	5151	2023	8847	3529
256	2609	562	4625	1240	7608	2156
512	2357	374	4245	703	7097	1346
1024	2236	311	4065	538	6615	924

We note that the parallel time on one computing node at Chinqchint cluster is equal to 3298 seconds, while the parallel time on one computing node at Graphene cluster is equal to 3115 seconds. The computing nodes at Chinqchint cluster compute slower than the computing nodes at Graphene cluster though they have twice as much computing cores due to the fact that the size of RAM memory at Graphene cluster is greater than the size of RAM memory at Chinqchint cluster and the parallel iterative methods perform frequent accesses to the memory. The more RAM memory, the more we avoid swapping and reduce the time to solve the problem.

Tables 5 and 6 display the computing times and computing gains, respectively, of synchronous, asynchronous, and hybrid schemes of computation for a different multi-cluster configuration. We consider two clusters that belong to the same site (Grenoble), ie, Edel and Genepi clusters of the Grid5000 testbed. The experiments are carried out with up to 32 computing nodes and a total of 256 computing cores (each computing node has 8 cores). Data exchange is made via Infiniband network in Edel and Genepi clusters and via Ethernet network (10 Gb/s) between them. We note that asynchronous iterations perform better than synchronous or hybrid iterations and that the combination of asynchronous iterations with GRIDHPC leads to important reduction in computing time.

Consider now multi-network configurations in Tables 3 and 5 and the respective computing times of asynchronous iterations for 128 computing cores, and then we see that the result in Table 5 (343 seconds) is less than in Table 3 (386 seconds). Nevertheless, the respective computing gain in Table 4 (8.06) is greater than the associated computing gain in Table 6 (5.72) since the parallel time on one computing node of Edel cluster

TABLE 8 Computing gain of parallel iterative methods applied to the obstacle problem with size 384^3 , 448^3 and 512^3 on Paravance cluster

Number of cores	Computing gain					
	384		448		512	
	Syn	Asyn	Syn	Asyn	Syn	Asyn
16	1	-	1	-	1	-
32	1.01	1.8	0.95	1.72	0.99	1.62
64	1.26	2.35	1.33	2.55	1.83	3.26
128	1.44	3.82	1.73	4.42	2.45	6.15
256	1.60	7.4	1.93	7.21	2.85	10.06
512	1.78	11.22	2.10	12.72	3.05	16.1
1024	1.87	13.5	2.2	16.62	3.28	23.49

is equal to 1962 seconds (computing nodes of Edel cluster are faster than computing nodes of Genepi cluster) while the parallel time on one computing node of Graphene cluster is equal to 3115 seconds.

We consider now a different set of the obstacle problems, ie, problems with size 384^3 , 448^3 and 512^3 . In this part of the study, we concentrate on task granularity and its effect on the computing gain of parallel iterative schemes of computation combined with GRIDHPC. The computing tests have been carried out on the Paravance cluster located in Rennes with 10 Gbps Ethernet network that has a large number of computing nodes and can be used to solve large instances of obstacle problem. We consider here up to 64 computing nodes and a total of 1024 computing cores. The results are displayed in Tables 7 and 8. The experiments show that asynchronous iterative schemes of computation achieve scalability when combined with GRIDHPC. The results show also that the computing gain generally increases when the problem size increases, ie, when the task granularity increases.

7 | CONCLUSIONS

In this paper, we have presented the GRIDHPC decentralized environment for high performance computing. GRIDHPC functionality relies on the reconfigurable communication protocol RMNP to support data exchange between computing nodes on multi-network systems with Ethernet, Infiniband, Myrinet, and on OpenMP for the exploitation of computing resources of multi-core CPU.

We have presented and analyzed a set of computational experiments with the decentralized environment GRIDHPC for a loosely synchronous application. In particular, we have studied the combination of GRIDHPC and parallel synchronous and asynchronous iterative schemes of computation for the obstacle problem in a multi-core and multi-network context. Our experiments are carried out on the Grid5000 platform with up to 1024 computing cores for a loosely synchronous application with frequent data exchange between computing nodes. We have considered several configurations like two Infiniband clusters connected via Ethernet or one Infiniband cluster and a Myrinet cluster connected via Ethernet. The results show that the combination of the GRIDHPC environment with asynchronous iterative algorithms scales up even when considering multi-cluster configurations.

In future work, we plan to use GRIDHPC for a different kind of loosely synchronous application. We shall consider the solution of large scale nonlinear network flow optimization problems. Finally, we shall extend GRIDHPC in order to combine multi-core CPUs and computing accelerators like GPUs.

ACKNOWLEDGMENTS

Part of this study has been made possible with funding of ANR-07-CIS7-011. Part of this work received also financial support of Government of Russian Federation under Grant 08-08. Experiments presented in this paper were carried out using the Grid5000 testbed, supported by a scientific interest group hosted by INRIA and including CNRS, RENATER, and several universities as well as other organizations (see <https://www.grid5000.fr>). The authors would also like to express their gratitude to the editor and the reviewers for their helpful remarks and comments.

ORCID

Bilal Fakih  <https://orcid.org/0000-0002-9520-4836>

REFERENCES

1. Hwang K, Fox G, Dongarra J. *Distributed and Cloud Computing: From Parallel Processing to the Internet of Things*. Burlington, MA: Morgan Kaufmann; 2012.
2. OpenMP. OpenMP: application programming interface. 2016. <http://www.openmp.org/wp-content/uploads/openmp-examples-4.5.0.pdf>

3. Wong GT, Hiltunen MA, Schlichting RD. A configurable and extensible transport protocol. In: Proceedings of the IEEE INFOCOM; 2001; Anchorage, AK.
4. Hiltunen MA, Schlichting RD. The cactus approach to building configurable middleware services. In: Proceedings of the Workshop on Dependable System Middleware and Group Communication (DSMGC); 2000; Nuremberg, Germany.
5. Aumage O, Bougé L, Denis A, et al. Madeleine II: a portable and efficient communication library for high-performance cluster computing. In: Proceedings IEEE International Conference on Cluster Computing (CLUSTER); 2000; Chemnitz, Germany.
6. Aumage O, Mercier G, Namyst R. MPICH/Madeleine: a true multi-protocol MPI for high performance networks. In: Proceedings 15th International Parallel and Distributed Processing Symposium (IPDPS); 2001; San Francisco, CA.
7. Heller T, Kaiser H, Schäfer A, Fey D. Using HPX and LibGeoDecomp for scaling HPC applications on heterogeneous supercomputers. In: Proceedings of the Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems (ScalA); 2013; Denver, CO.
8. Kaiser H, Heller T, Berge A, Adelstein-Lelbach B. HPX V0.9.10: a general purpose C++ runtime system for parallel and distributed applications of any scale. 2015. <http://github.com/STELLAR-GROUP/hpx>
9. Kaiser H, Brodowicz M, Sterling T. ParallelX: an advanced parallel execution model for scaling-impaired applications. Paper presented at: 2009 International Conference on Parallel Processing Workshops; 2009; Vienna, Austria.
10. Stumpf M. Distributed GPGPU Computing with HPXCL. 2014. Paper presented at: LA-sigma TESC Meeting; 2014; Baton Rouge, LA.
11. Huck K, Shende S, Malony A, et al. An early prototype of an autonomic performance environment for exascale. In: Proceedings of the 3rd International Workshop on Runtime and Operating Systems for Supercomputers (ROSS); 2013; Eugene, OR.
12. Augonnet C, Thibault S, Namyst R, Wacrenier P-A. StarPU: a unified platform for task scheduling on heterogeneous multicore architectures. In: Proceedings of the 15th European Conference on Parallel Processing; 2009; Delft, The Netherlands.
13. Fakihi B, El Baz D. Heterogeneous computing and multi-clustering support via peer-to-peer HPC. Paper presented at: 2018 26th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP); 2018; Cambridge, UK.
14. Nguyen TT, El Baz D, Spitéri P, Jourjon G, Chau M. High performance peer-to-peer distributed computing with application to obstacle problem. In: Proceedings of the 24th 2010 IEEE International Symposium on Parallel and Distributed Processing, Workshops and Phd Forum (IPDPSW); 2010; Atlanta, GA.
15. Nguyen TT, El Baz D, Spitéri P, Garcia T, Chau M. Asynchronous peer-to-peer distributed computing for financial applications. Paper presented at: 2011 IEEE International Symposium on Parallel and Distributed Processing Workshops and Phd Forum (IPDPSW); 2011; Shanghai, China.
16. Chazan D, Miranker W. Chaotic relaxation. *Linear Algebra Appl.* 1969;2:199-222.
17. Baudet GM. Asynchronous iterative methods for multiprocessors. *J Assoc Comput Mach.* 1978;25:226-244.
18. El Baz D, Frommer A, Spitéri P. Asynchronous iterations with flexible communication: contracting operators. *J Comput Appl Math.* 2005;176:91-103.
19. El Baz D. An efficient termination method for asynchronous iterative algorithms on message passing architectures. In: Proceedings of the International Conference on Parallel and Distributed Computing Systems; 1996; Dijon, France.
20. Grid5000 platform. <http://www.grid5000.fr>.

How to cite this article: Fakihi B, El Baz D, Kotenko I. GRIDHPC: A decentralized environment for high performance computing. *Concurrency Computat Pract Exper.* 2019;e5320. <https://doi.org/10.1002/cpe.5320>