Assessing the Performance of MPI Applications Through Time-Independent Trace Replay

Frédéric Desprez, George Markomanolis, Martin Quinson, Frédéric Suter

► To cite this version:

Frédéric Desprez, George Markomanolis, Martin Quinson, Frédéric Suter. Assessing the Performance of MPI Applications Through Time-Independent Trace Replay. Second International Workshop on Parallel Software Tools and Tool Infrastructures (PSTI 2011) Held in conjunction with ICPP 2011, the 40th International Conference on Parallel Processing, Sep 2011, Taipei, Taiwan. inria-00546992

HAL Id: inria-00546992

<https://hal.inria.fr/inria-00546992>

Submitted on 15 Dec 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Assessing the Performance of MPI Applications Through Time-Independent Trace Replay

Frédéric Desprez, George S. Markomanolis (INRIA, LIP, ENS Lyon, Lyon, France)

Martin Quinson (Nancy University, LORIA, INRIA, Nancy, France)

Frédéric Suter (Computing Center, CNRS, IN2P3, Lyon-Villeurbanne, France)

N° 7489

2010

Domaine 3

**apport
de recherche**

Assessing the Performance of MPI Applications Through Time-Independent Trace Replay

Frédéric Desprez, George S. Markomanolis (INRIA, LIP, ENS Lyon,
Lyon, France)

Martin Quinson (Nancy University, LORIA, INRIA, Nancy, France)
Frédéric Suter (Computing Center, CNRS, IN2P3, Lyon-Villeurbanne,
France)

Domaine : Réseaux, systèmes et services, calcul distribué
Équipe-Projet AlGorille, Graal

Rapport de recherche n° 7489 — 2010 — 21 pages

Abstract: Simulation is a popular approach to obtain objective performance indicators platforms that are not at one's disposal. It may help the dimensioning of compute clusters in large computing centers. In this work we present a framework for the off-line simulation of MPI applications. Its main originality with regard to the literature is to rely on time-independent execution traces. This allows us to completely decouple the acquisition process from the actual replay of the traces in a simulation context. Then we are able to acquire traces for large application instances without being limited to an execution on a single compute cluster. Finally our framework is built on top of a scalable, fast, and validated simulation kernel.

In this paper, we present the used time-independent trace format, investigate several acquisition strategies, detail the developed trace replay tool, and assess the quality of our simulation framework in terms of accuracy, acquisition time, simulation time, and trace size.

Key-words: Message Passing Interface; Off-line simulation; Performance prediction.

Rejeu de traces indépendantes du temps pour l'étude des performances d'applications MPI

Résumé : La simulation est une approche très populaire pour obtenir des indicateurs de performances objectifs sur des plates-formes qui ne sont pas disponibles. Cela peut permettre le dimensionnement de grappes de calculs au sein de grands centres de calcul. Dans cet article nous présentons un outil de simulation *post-mortem* d'applications MPI. Sa principale originalité au regard de la littérature est d'utiliser des traces d'exécution indépendantes du temps. Cela permet de découpler intégralement le processus d'acquisition des traces de celui de rejeu dans un contexte de simulation. Il est ainsi possible d'obtenir des traces pour de grandes instances de problèmes sans être limité à des exécutions au sein d'une unique grappe. Enfin notre outil est développé au dessus d'un noyau de simulation scalable, rapide et validé.

Cet article présente le format de traces indépendantes du temps utilisé, étudie plusieurs stratégies d'acquisition, détaille l'outil de rejeu que nous avons développé, et évalué la qualité de nos simulations en termes de précision, temps d'acquisition, temps de simulation et tailles de traces.

Mots-clés : MPI, simulation hors ligne, prédiction de performances.

1 Introduction

Computational Science is the third scientific way to study problems arising in various domains such as physics, biology, or chemistry. It is complementary to theory and actual experiments and consists in conducting studies *in silico*. This approach makes an heavy use of resources located in computing centers. As the number of scientific domains producing results from *in silico* studies increases, the computing centers then have to upgrade their infrastructures in a continuous way. The resources impacted by such upgrades are of different kinds, e.g., computing, network and storage. Moreover each kind of resource grows more complex with each generation. Processors have more and more cores, low latency and high bandwidth network solutions become mainstream and some disks now have access time close to that of memory.

The complexity of the decision process leading to the evolution of a computing center is then increasing. This process often relies on years of experience of system administrators and users. The former knows how complex systems work while the latter have expertise on the behavior of their applications. Nevertheless this process lacks of objective data about the performance of a given candidate infrastructure. Such information can only be obtained once the resources have been bought and the applications can be tested. Any unforeseen behavior can then lead to tremendous but vain expenses.

Many simulation frameworks have been proposed over the last decade to obtain objective indicators beforehand. Most of them focus on parallel applications relying on the Message Passing Interface (MPI) [1]. The simulation approaches used in these frameworks fall into two categories: *off-line simulation*, also called trace-based simulation or post-mortem simulation, and *on-line simulation*, also called simulation via direct execution. In off-line simulation a log of a previous execution of the application is “replayed” on a simulated platform. In on-line simulation the application is executed but part of the execution takes place within a simulation component.

In this work we present a framework to simulate MPI applications following the off-line approach. This effort is part of the SIMGrid project [2]. It leverages the simulation existing techniques and models in SIMGrid. More specifically, this work makes these main contributions with respect to the MPI simulation literature:

1. Propose a new execution log format that is independent of time; This format includes volumes of computation and communications for each event instead of classical time-stamps;
2. An original approach that totally decouples the acquisition of the trace from its replay; Several scenarios can then be proposed that allow for the acquisition of large execution traces;
3. A trace replay tool on top of fast, scalable and validated simulation kernel; This ensures the efficiency and quality of our off-line simulation framework;
4. Experimental results that show the simulation accuracy, the acquisition time, the simulation time, and the trace size, of our off-line simulation framework.

This paper is organized as follows. Section 2 reviews related work. Section 3 introduces the format of a time-independent trace of an MPI application while Section 4 details the acquisition process of such traces. Section 5 explain how time-independent traces can be replayed within the SIMGrid simulation framework. Our prototype is evaluated in Section 6 with regard to different metrics. Section 7 concludes with a summary of results and perspectives on future research directions.

2 Related Work

One option for simulating the execution of an MPI application is *on-line simulation*. In this approach, the actual code, with no or only marginal modification, is executed on a *host platform* that attempts to mimic the behavior of a *target platform*, i.e., a platform with hardware characteristics different from those of the platform on which traces were obtained. Part of the instruction stream is then intercepted and passed to a simulator. LAPSE is a well-known on-line simulator developed in the early 90's [3] (see therein for references to precursor projects). In LAPSE, the parallel application executes normally but when a communication operation is performed a corresponding communication delay is simulated on the target platform using a simple network model (affine point-to-point communication delay based on link latency and bandwidth). MPI-SIM [4] builds on the same general principles, with the inclusion of I/O subsystem simulation in addition to network simulation. A difference with LAPSE is that MPI processes run as threads, a feature which is enabled by a source code pre-processor. Another project similar in intent and approach is the simulator described in [5]. The BigSim project [6] also builds on similar ideas. However, unlike MPI-SIM, BigSim allows the simulation of computational delays on the target platform. This makes it possible to simulate "what if?" scenarios not only for the network but also for the compute nodes of the target platform. Simulation of computation delays in BigSim is done based either on user-supplied projections for the execution time of each block of code, or on scaling execution times measured on the host platform by a factor that accounts for the performance differential between the host and the target platform, or based on sophisticated execution time prediction techniques such as those developed in [7]. The weakness of this approach is that since the computational application code is not executed, the computed application data is erroneous. Consequently, application behavior that is data-dependent is lost. This is acceptable for many regular parallel applications, but can make the simulation of irregular applications (e.g., branch-and-bound or sparse matrices computations) questionable at best. Going further, the work in [8] uses a cycle-accurate hardware simulator of the target platform to simulate computation delays, which leads to a high ratio of simulation time to simulated time.

One difficulty faced by all above MPI-specific on-line simulators is that the simulation, because done through a direct execution of the MPI application, is inherently distributed. Parallel discrete event simulation raises difficult correctness issues pertaining to process synchronization. For the simulation of parallel applications, techniques have been developed to speed up the simulation while preserving correctness (e.g., the asynchronous conservative simulation algorithms in [9], the optimistic simulation protocol in [6]). A solution could be to run the simulation on a single node but it requires large amounts of CPU and RAM resources. For most aforementioned on-line approaches, the resources required to run a simulation of an MPI application are commensurate to those of that application. In some cases, those needs can even be higher [8, 10]. One way to reduce the CPU needs of the simulation is to avoid executing computational portions of the application and simulate only expected delays on the target platform [6]. Reducing the need for RAM resources is more difficult and, in general, if the target platform is a large cluster, then the host platform must also be a large cluster.

One approach that avoids these particular challenges, but that comes with challenges of its own, is *off-line simulation*. In this approach, which we use in this work, a log, or trace, of MPI communication events (time-stamp, source, destination, data size) is first obtained by running the application on a real-world platform. A simulator then replays the execution of the application as if it were running on a *target platform*. This

approach has been used extensively, as shown by the number of trace-based simulators described in the literature since 2009 [11, 12, 13, 14, 15]. The typical approach is to compute the durations of the time intervals between MPI communication operations, or “CPU bursts.” Upon replaying the application, the CPU bursts are modified to account for the performance differential between the platform used to obtain the trace and the target platform, either using simple scaling [12, 13, 15] or using a more sophisticated convolution between the application computational signature and the target platform’s hardware signature [7]. Network communications are simulated based on the communication events recorded in the trace and on a simulation model of the network. This tight link between *host* and *target* platforms, caused by the time-stamps in the traces, is a potential drawback of the off-line approach. Indeed it limits the usage of the traces to similar platform for which simple scaling is possible.

A challenge for off-line simulation is the large size of the traces, which can prevent running the simulation on a single node. Mechanisms have been proposed to improve scalability, including compact trace representations [12] and replay of a judiciously selected subset of the traces [14]. Another challenge is that it is typically necessary to obtain the trace on a platform that has the same scale as the target platform. However, trace extrapolation to larger numbers of nodes than that of the platform used to obtain the trace is feasible in some cases [11, 13].

For both approaches, the complexity of the network simulation model has a high impact on speed and scalability, thus compelling many authors to adopt simplistic network models. One simplification, for instance, is to use monolithic performance models of collective communications rather than simulating them as sets of point-to-point communications [12, 16]. Another simplification used in most aforementioned simulators, whether off-line or on-line, is to ignore network contention because it is known to be costly to simulate [17]. The work in [12] proposes the use of simple analytical models of network contention for off-line simulation. An exception is the MPI-NetSim on-line simulator [10], which provides full-fledge contention simulation via a packet-level discrete event network simulator. As a result, the simulator may run more slowly than the application.

3 Time-Independent Trace Format

All the off-line MPI simulators reviewed in the previous section relies on timed traces, i.e., each occurring event is associated to a time-stamp. Consequently a simulator has to apply a correction factor to these time-stamps when it simulates the execution on the target platform. This implies to know precisely what are the respective performance of the host and target platforms. In other words, each execution trace must come with an accurate description of how it has been acquired. Finally determining the right scaling factor can be tedious depending on the degree of similarity of both platforms.

To free ourselves of these constraints related to time-stamps, we propose in this work to produce *time-independent* traces. For each event occurring during the execution of the application, e.g., a CPU burst or a communication operation, we log the volume of the operation (in number of floating operations or bytes) instead of the time spent to execute it. Indeed this type of information does not vary with the characteristics of the host platform. For instance, the size of the messages sent by an application is not likely to change according to the specifics of the network interconnect, while the computation amount performed within a `for` loop does not increase with the processing speed of a CPU. This claim is not valid for adaptive MPI applications that modify

their execution path according to the execution platform. This type of applications, that represents only a small fraction of all MPI applications, is not covered by our work.

A time-independent trace can then be seen as a list *actions*, e.g., computations and communications, performed by each process of an MPI application. An action is described by the *id* of the process that does this action, a *type*, e.g., a computation or a communication operation, a *volume*, i.e., a number of flops or bytes, and some action specific parameters, e.g., the id of the receiving process for a one-way communication.

Figure 1: MPI code sample of some computation on a ring of processes (left) and its equivalent time-independent trace (right).

The left hand side of Figure 1 shows a simple computation on a ring of four processes. Each process computes one million floating points operations and send one million bytes to its neighbor. The right hand side of this figure displays the corresponding time-independent trace. Note that, depending on the number of processes and the number of actions, it may be preferable to split the time-independent trace in several files, e.g., one file per process.

MPI actions	Trace entry
CPU burst	<id> compute <volume>
MPI_Send	<id> send <dst_id> <volume>
MPI_Isend	<id> Isend <dst_id> <volume>
MPI_Recv	<id> recv <src_id> <volume>
MPI_Irecv	<id> Irecv <src_id> <volume>
MPI_Broadcast	<id> bcast <volume>
MPI_Reduce	<id> reduce <vcomm> <vcomp>
MPI_Allreduce	<id> allReduce <vcomm> <vcomp>
MPI_Barrier	<id> barrier
MPI_Comm_size	<id> comm_size #proc
MPI_Wait	<id> wait

Table 1: time-independent counterparts of the actions performed by each process involved in a MPI application.

Table 1 lists all the MPI functions for which there is a corresponding action implemented in our first prototype. For the collective operations, we consider that all the processes are involved as the `MPI_Comm_split` function is not implemented. Another design choice is to root these collective operations on process 0. Finally the `comm_size` action has to appear in the trace file associated to each process prior to any collective operation.

4 Trace Acquisition Process

In this section we detail the acquisition process of a time-independent execution trace that can be replayed within a simulation environment. This process, depicted in Figure 2 comprises four steps: (i) the instrumentation of the application; (ii) the execution of the instrumented version of the application; (iii) the extraction of the action list for each process; and (iv) the gathering of the different traces into a single node. In what follows we give some details on each of these steps.

Figure 2: Time-independent traces acquisition process.

4.1 Instrumentation

The first step of the acquisition process is to instrument the application. We base our prototype on TAU [18] for its non-intrusive instrumentation method. TAU is actually a profiling tool that offers a tracing features. To enable them the `-TRACE` flag has to be used and some environment variables, e.g., `TAU_TRACK_MESSAGE` to track message sender and receiver ids, have to be set. The semi-automatic instrumentation described in what follows is activated by the `-pdt` flag, while the access to hardware counters through the PAPI [19] interface is allowed by adding the `-papi` flag to the command line.

An interesting feature of TAU is selective instrumentation. It can be done in different ways. One consists in listing in a separate file which functions have (or have not) to be traced. All the functions in the call path of the listed functions will also be traced. However, this technique may not be enough to isolate a given call, e.g., if a function is called twice but only one call has to be traced. A solution is then to insert two macros provided by the tau API, namely `TAU_ENABLE_INSTRUMENTATION` and `TAU_DISABLE_INSTRUMENTATION` in the source code.

The following example illustrates the instrumentation of the `SSOR(itmax)` function call in a LU factorization.

```

1      call TAU_ENABLE_INSTRUMENTATION()
2      call ssor(itmax)
3      call TAU_DISABLE_INSTRUMENTATION()

```

One additional call to each of these macros is required to define neat disable/enable sections. Such slight modifications of the source code can be handled by the pre-processor. Indeed the initial program has to be compiled using one of the scripts provided by TAU, i.e., `tau_cc.sh` and `tau_f90.sh` for C and Fortran codes respectively.

4.2 Execution

As shown in Figure 2, a time-independent trace can be acquired in many different ways. The only mandatory parameter is the number of participating processes, as the trace only comprises information about computation and communication volumes. We remind that the off-line approach is not suited for irregular and adaptive applications and that regular applications represent a large part of current MPI codes used in production.

Figure 2 shows three possible acquisition methods that consist in running the MPI application in a:

- **Regular mode:** with one process per CPU. This is the way the other off-line simulators obtain execution traces. Indeed, such an acquisition prevents abnormal timings due to resource access concurrency. The main drawback of this acquisition mode is to require as many processors as comprised in the target platform to get a trace. Its scalability is thus limited and the acquisition of traces for larger is not possible.
- **Folding mode:** with more than one process per CPU. This allows for the acquisition of traces for larger instances of the application or to use less resources. The folding factor is obviously limited by the available amount of memory on the involved computing nodes.
- **Scattering mode:** where the CPUs do not necessarily belong to the same commodity cluster. More nodes than what is available on the target platform can thus be used to acquire traces of large instance. This mode thus tackles the issues of the regular mode.

A fourth acquisition can also be envisioned. The **Scattering and Folding** mode is, as the name says, the combination of the last two modes. It further increases the scalability of the acquisition process. Apart from the *regular* mode, all the proposed acquisition methods are possible only because the independence to time of the produced traces. Consequently, other off-line simulators can not take benefits of these alternate acquisition modes.

4.3 Post-processing of the Execution Traces

When the execution of a program instrumented with TAU completes, many files are produced. They fall in two categories: *trace* files and *event* files. The generated trace files are named:

```
tautrace.<node>.<context>.<thread>.trc,
```

where `<node>` is the rank of the MPI process whose execution is logged in the file. The two other fields, i.e., `<context>` and `<thread>`, are only used for multi-threaded applications. In this case, TAU distinguishes each thread and groups the threads according to the virtual address space they share.

A *trace file* is a binary file that includes all the events that occur during the execution of the application for a given process. For each event, this file indicates when this event (e.g., a function call or an instrumented block) starts and finishes. The time spent and the number of computed floating point operations between these begin/end tags are also stored. For MPI events all the parameters of the MPI call, including source, destination, and message size, are stored.

To reduce the size of the trace files, TAU stores a unique id for each traced event instead of its complete signature. The matching between the ids and the the functions descriptions can be found in the *event files*. These files are named:

`events.<node>.edf.`

There is only one event file per MPI process. Each event file contains information about each traced function. For any function, an event file stores its numerical id, the group it belongs to, e.g., `MPI` for all MPI functions, a tag to distinguish TAU events from those defined by the user, and the *name type* which is the actual name of the traced function. Some extra parameters required by TAU can also be stored into an event file. For instance, the keyword `EntryExit` is used to declare a function that occurs between two separate events, i.e., entry and exit. Conversely the `TriggerValue` keyword typically corresponds to a counter that increases monotonically from the beginning of the execution. Such a trigger has to be activated twice to determine the evolution of the counter value during the corresponding period of time.

The following example shows two entries of an event file generated by TAU that corresponds respectively to the `MPI_Send` function and to the access to an hardware counter that measures the number of floating point operations.

```
49 MPI 0 "MPI_Send()" " EntryExit
1 TAUevent 1 "PAPI_FP_OPS" TriggerValue
```

To perform an off-line simulation of a MPI application with SIMGrid, two steps are mandatory. First we have to *extract* a time-independent trace from the trace and event files produced by TAU. Second we have to *gather*, and sometimes *merge*, the extracted traces on a single node where the replay takes place.

As the trace files generated by TAU are binary files, there is a need for an interface to extract information. Such an API is provided by the TAU Trace Format Reader library (TFR) [20]. This tool provides the necessary functions to handle a trace file, including a function to read events. It also defines a set of eleven callback methods, that correspond to the different kinds of events that appear in a TAU trace file. For instance there are callbacks for entering or exiting a function and triggering a counter. The implementation of these callback methods is let to the developer.

We thus developed a C/MPI parallel application, called `tau2simgrid`, that implements the different callback methods of the TFR library. This program basically opens, in parallel, all the TAU trace files and read them line by line. For each event, the corresponding callback function is called. To illustrate how `tau2simgrid` extracts the necessary data to produce a time-independent trace, we detail the case of a call to the `MPI_Send` function. Figure 3 presents the parameters of the different callbacks related to this function call on process 1 in a readable format. Each line starts by the process id, the thread id, the time at which the event occurred and the name of the event. The remaining fields are event dependent.

As mentioned earlier, the event that corresponds to a `MPI_Send` is tagged as `EntryExit` in the event file with the event id 49. The first occurring callback will then be on the `EnterState` function (line 1). The matching `LeaveState` event (line 6)

1	1	0	1.42947e+06	EnterState	49				
2	1	0	1.42947e+06	EventTrigger	1	164035532			
3	1	0	1.4295e+06	EventTrigger	46	163840			
4	1	0	1.4295e+06	SendMessage	0	0	163840	1	0
5	1	0	1.4299e+06	EventTrigger	1	164035624			
6	1	0	1.4299e+06	LeaveState	49				

Figure 3: List of callbacks related to a call to the `MPI_Send` function.

define the scope of events related to the function call. Four events are enclosed between these boundaries. Two of them (lines 2 and 5) correspond to the hardware counter measuring the number of flops, as identified in the event file and indicated by the event id 1 in the trace file. These two events are used to respectively ends the CPU burst preceding the MPI call and starts the next one. The number of flops computed within a MPI call, mainly due to buffer allocation costs, are ignored as they are accounted for by the network model. The last two events are related to the sent message. The `EventTrigger` on line 3 only provides the size of the message (163,840 bytes), which is not enough to build an entry in the time-independent trace. The `SendMessage` event (line 4) gives more information, namely the process and thread ids of the receiver, the size of the message, and the MPI tag and communicator for this communication.

Thanks to all these information extracted from both TAU trace and event files, our `tau2simgrid` application can generate the following entry of a time-independent trace:

```
p1 send p0 163840
```

Note that for asynchronous and collective communications, the extraction process is more complex. For instance, the mandatory information to write the entry corresponding to a `MPI_Irecv`, e.g., the receiver id, are given by the `RecvMessage` event which generally occurs within the `MPI_wait` function. This implies to implement some lookup techniques in `tau2simgrid` to retrieve all the necessary parameters.

After the execution of `tau2simgrid`, the produced traces can be injected into the simulation framework. However, they first have to be gathered on a single node onto which the simulation will take place. File gathering is a problem that has been studied for a long time. A common and efficient approach is to rely on a K-nomial tree reduction allowing for $\log_{(K+1)} N$ steps, where N is the total number of files, and K is the arity of the tree. We developed a simple script to perform such a gathering. This script can be configured to adapt the arity to the total number of traces and the number of computed nodes involved in the trace acquisition.

The cost, in terms of execution time, of `tau2simgrid` and the gathering script will be assessed in Section 6.

5 Trace Replay with SimGrid

Our framework to simulate MPI applications following the off-line approach is tightly connected to the `SIMGrid` project [2]. `SIMGrid` provides core functionalities for the simulation of distributed applications in heterogeneous distributed environments. The main goal of `SIMGrid` is to facilitate the research in the area of parallel and distributed

large scale systems such as Grids, P2P systems and clouds. SIMGrid relies on a scalable and extensible simulation engine and offers several user APIs.

Since release 3.3.3, SIMGrid allows users to describe an applicative workload as a time-independent trace such as that described in Section 3. Figure 4 shows the different components that are needed to replay such traces with SIMGrid. Apart from the *time-independent trace(s)*, a description of the *target platform* and the *deployment* of the application, i.e., how processes are mapped onto processors, are also passed to the *trace replay tool* which, in turn, is built on top of the *simulation kernel*. Decoupling the simulation kernel, and then the simulator, from the simulation scenario offers a greater flexibility. This way a wide range of “what if?” scenarios can be explored without any modification of the simulator. Changing the input files of the *trace replay tool* is enough.

Figure 4: Inputs and outputs of the SIMGrid trace replay framework.

The *simulation kernel* of SIMGrid relies on macroscopic models for computation resources. Tasks costs are expressed in number of floating point operations (flops). The computing power of a CPU is then in flop/s. For network resources, SIMGrid uses an analytical network contention model. This model was developed for arbitrary network topologies with end-points that use standard network protocols, such as TCP/IP, and are connected via multi-hop paths. Instead of being packet-based, the model is flow-based, meaning that at each instant the bandwidth allocated to an active flow (i.e., a data transfer occurring between two end-points) is computed analytically given the topology of the network and all currently active flows. This model is described and validated via comparison to the GTNetS packet-level simulator in [21]. While this generic model is applicable to networks ranging from local-area to wide-area networks, it can be specialized for cluster interconnects. An original model has been recently added to the SIMGrid simulation kernel to take the specifics of MPI implementations on compute cluster interconnects using TCP into account. For instance, a message under 1 KiB fits within an IP frame, in which case the achieved data transfer rate is higher than for larger messages. Also, MPI implementations for `MPI_Send()` typically switch from buffered to synchronous mode above a certain message size. Consequently, instead of being an affine function of message size, communication time is *piece-wise linear*. In practice, this model is instantiated for 3 segments, leading to 8 parameters defining the model (2 for defining the boundaries of the 3 segments, and one latency and bandwidth parameter for each segment).

Figures 5 and 6 respectively present a *platform file* and a *deployment file* that correspond to the scenario of Figure 1. The platform described here is a compute cluster that comprises four homogeneous machines interconnected through a switched network. The deployment file indicates on which node of the cluster each process will run. For instance, the MPI process of rank 0 will be executed on the node named `mycluster-0.mysite.fr`

```
<?xml version='1.0'?>
<!DOCTYPE platform SYSTEM "simgrid.dtd">
<platform version="3">
  <AS id="AS_mysite" routing="Full">
    <cluster id="AS_mycluster"
      prefix="mycluster-" suffix=".mysite.fr"
      radical="0-3" power="1.17E9"
      bw="1.25E8" lat="16.67E-6"
      bb_bw="1.25E9" bb_lat="16.67E-6"/>
  </AS>
</platform>
```

Figure 5: Example of SIMGrid platform file corresponding to the MPI application described in Figure 1.

```
<?xml version='1.0'?>
<!DOCTYPE platform SYSTEM "simgrid.dtd">
<platform version="3">
  <process host="mycluster-0.mysite.fr"
    function="p0"/>
  <process host="mycluster-1.mysite.fr"
    function="p1"/>
  <process host="mycluster-2.mysite.fr"
    function="p2"/>
  <process host="mycluster-3.mysite.fr"
    function="p3"/>
</platform>
```

Figure 6: Example of SIMGrid deployment file corresponding to the MPI application described in Figure 1.

Figure 4 also shows the possible outputs of an off-line simulation. In this work, we focus on obtaining a *simulated execution trace*. This kind of output provides an estimation of the execution time of the target application in the particular experimental scenario described by the platform and deployment files. It is also possible to generate a timed trace that corresponds to this particular scenario by adding timers (measuring simulated time) in the trace replay tool. Finally it would also be interesting to derive a profile of the application from this timed trace. But this last kind of output requires complex analysis tools such as those develop in the TAU and Scalasca [22] projects.

To replay a time-independent trace with SIMGrid, a simulator, the *trace replay tool* has to be written on top of the simulation kernel using one of the APIs provided by SIMGrid. Our prototype uses the MSG API. This simulator has to:

1. Include a function that corresponds to the expected behavior of a given action.
This has to be done for each action that occurs in the trace. Generally such

function just calls one or several MSG functions. For instance, the code for the `compute` action is

```

1 static void compute(xbt_dynar_t action){
2     char *amount = xbt_dynar_get_as(action,
3                                     2, char *);
4     m_task_t task = MSG_task_create(NULL,
5                                     parse_double(amount),
6                                     0, NULL);
7     MSG_task_execute(task);
8     MSG_task_destroy(task);
9 }

```

An entry of the trace file is passed to this function as a dynamic array (`xbt_dynar_t`) of strings, one for each field of the entry. Once the amount of flops to compute has been extracted (lines 2-3), it is possible to create (lines 4-6) the corresponding SIMGrid task and execute it (line 7). This SIMGrid task is destroyed (line 8) as soon as the computation of `amount` flops has been simulated.

2. Register this function with `MSG_action_register`. For instance to link the `compute` keyword in the trace to the above function, the main function of the simulator includes the following call

```

1 MSG_action_register("compute", compute);

```

3. Call the function `MSG_action_trace_run` that takes either a trace file name or `NULL` as input. When no file name is given, this mean that there exists one trace file per process. In this case, the names of these trace files are given in the platform file, as shown below.

```

1 <process host="mycluster-1.mysite.fr"
2     function="pl">
3     <argument value="SG_process1.trace"/>
4 </process>

```

An essential step to make accurate performance predictions through trace replay is the calibration of the simulation framework. In the case of SIMGrid, it consists in instantiating the platform file with pertinent values. In other words the number of floating operations a CPU can compute in one second and the latency and bandwidth of communication links have to be set. Such a calibration strongly depends on both application and execution environment. Indeed different types of computation may lead to different flop rates on a given CPU. This is mainly due to how efficiently the computation can use the different levels of cache. Moreover the performance of a given computation may differ with regard to the processor brand.

We instantiate the flop rate of each host in the platforms as follows. A small instrumented instance of the target application is run of the platform to describe. This allows us to determine the number of flops of each event as long as the time spent to computed them. Then we can determine a flop rate of each single action, compute a weighted average on each process, and get an average flop rate for all the process set. Finally we repeat this this procedure five times and compute an average over these five runs to smoothe the runtime variations. We use this final value to instantiate the SIMGrid platform file.

The instantiation of the network parameters of the platform file is done in two steps. To set the bandwidth, we use the nominal value of the links, e.g., 1 GiB for GigaEthernet links. For the latency of a communication link, we rely on the `Pingpong_Send_Recv` experiment of the SKaMPI [23] benchmark suite. We take the value obtained for a 1-byte message and divided it by six. This factor of six comes from two sources. We have to divide the ping-pong time by two to obtain the latency of a one-way message. Then we divide it by three to take the topology of a cluster into account. Indeed, two nodes in a compute cluster are generally connected through two links and one switch. In case of hierarchical network, we account for this hierarchy in the determination of the latency.

The second step consists in instantiating the piece-wise linear model used by SIMGrid dedicated to MPI communications on compute clusters. SIMGrid provides a Python script that takes as input the latency and bandwidth determined as above, the output of the SKaMPI run, and the number of links connecting the two nodes used for the ping-pong. Then this script determines the latency and bandwidth correction factors that lead to a best-fit of the experimental data for each segment of this piece-wise linear model.

6 Experimental Evaluation

6.1 Experimental Setup

To conduct this evaluation, we use the LU factorization (LU) program, which is part of the NAS Parallel Benchmarks (NPB) suite. The NPB are a set of programs commonly used to assess the performance of parallel platforms. Each benchmark can be executed for 7 different *classes*, denoting different problem sizes: S (the smallest), W, A, B, C, D, and E (the largest). For instance, a class D instance corresponds to approximately 20 times as much work and a data set almost 16 as large as a class C problem. We selected the LU factorization because it mixes computations and communications and is a building block of many scientific applications.

To acquire trace-independent traces we used two clusters of the Grid’5000 experimental platform: *bordereau* and *gdx*. The *bordereau* cluster comprises 93 2.6GHz Dual-Proc, Dual-Core AMD Opteron 2218 nodes. All these nodes are connected to a single 10 Gigabit switch. The *gdx* cluster comprises 186 2.0 GHz Dual-Proc AMD Opteron 246 scattered across 18 cabinets. Two cabinets share a common switch and all these switches are connected to a single second level switch through Ethernet 1 Gigabit links. Consequently a communication between two nodes located in two distant cabinets goes through three different switches. These two clusters are interconnected through a dedicated 10 Gigabit network.

	Acquisition mode	R	F-2	F-4	F-8	F-16	F-32	S-2	SF-(2,2)	SF-(2,4)	SF-(2,8)	SF-(2,16)
	Number of nodes	64	32	16	8	4	2	(32,32)	(16,16)	(8,8)	(4,4)	(2,2)
B	Execution Time (in sec.)	20.73	52.96	88.66	179.07	347.27	689.18	37.54	79.19	134.05	277.25	505.64
	Ratio to regular mode	1	2.55	4.28	8.64	16.75	33.25	1.81	3.82	6.47	13.37	24.39
C	Execution Time (in sec.)	57.77	143.45	272.45	511.75	1,011.59	1,970.05	85.71	211.95	421.71	772.56	1,442.79
	Ratio to regular mode	1	2.22	4.13	7.79	15.14	31.79	1.48	3.67	7.3	13.37	24.97

Table 2: Evolution of the execution time of an instrumented LU benchmark executed by 64 processes with regard to the acquisition mode. Results obtained on the *bordereau* and *gdx* clusters using one core per node.

One of the key concept of the Grid'5000 experimental platform is to offer its users the capacity to deploy their own system image at will. For our experiments we built a Debian Lenny image. The kernel (v2.6.25.9) was patched with the `perfctr` driver (v2.6.38) to enable hardware counters. We also installed TAU (v2.18.3) and its software dependencies on Program Database Toolkit (v3.14.1), which is used for automatic instrumentation, and PAPI (v3.7.0), that provides access to the hardware counters. We built the NAS Parallel Benchmarks (v3.3) on top of OpenMPI (v1.3.3). Finally the traces are replayed in SIMGrid (v3.6-r9069).

6.2 Evaluation of the Acquisition Modes

The main characteristic of the proposed acquisition process is to be totally decoupled from the target platform. As mentioned in Section 4, several modes can be used to acquire a time-independent trace. The first experiment of this section investigates the distribution of the acquisition time among the different steps, i.e., execution, instrumentation, extraction, and gathering. Figure 7 shows such a distribution for the acquisition of time-independent traces of the LU benchmark for classes B and C and for different number of processes. These results were obtained on the *bordereau* cluster following the **Regular** acquisition mode and deploying only one process per node. We run the complete acquisition process ten times and show the average value for each part of it.

Figure 7: Distribution of the acquisition time for different instances of the LU benchmark in the *regular* acquisition mode.

As more processes are involved in the computation of the LU factorization, the time needed to run the application, get a TAU trace, and extract the time-independent information decreases linearly with the number of processes. This direct benefit of parallelism exploitation shows its limits when the sequential part of the execution becomes too small, as for Class B on 64 processes, for instance. Conversely, the time needed to gather all the generated traces on a single node, with a 4-nomial tree, increases with the number of processes as the depth of the reduction tree also grows. However this step remains the least time consuming.

If we focus on the part of the acquisition time that is strictly related to the production of time-independent traces, i.e., the extraction and gathering steps, we see that it represent at most 34.91% of the total acquisition time. The worst value is obtained for the Class B with 64 processes instance for which the execution time of the benchmark is the smallest. However, large number of processes are generally used to solve large problem instances. Then we can conclude that the extra overhead required to get a time-independent trace can be afforded. Moreover such traces need to be acquired only once and can then be used to explore much more “what if?” scenarios than with timed traces.

Now we estimate the impact of the **Folding** (F- x), **Scattering** (S- y), and **Scattering and Folding** (SF-(u, v)) acquisition modes on the execution time of the instrumented version of the LU benchmark when compared to the **Regular** mode (R). When processes are folded on a smaller number of nodes, x denotes the folding factor. For instance F-4 means that four processes are executed on a single CPU. In the scattering mode, y is the number of sites used during the acquisition. Finally when both modes are combined, SF-(u, v) means that the execution is scattered over u sites and that v processes run on each node. Table 2 presents the results of this comparison in terms of execution time and performance degradation. In all these experiments, we use only one core per node.

We see that the time needed to execute the instrumented application increases roughly linearly with the folding factor (F-2 to F-32). This was expected as the codes of several processes have to be executed concurrently on a single CPU. However, there is no extra overhead induced by this process folding approach. When the execution is scattered across two different clusters (S-2), the overhead comes from two distinct factors. First, some communications are now made on a wide area network and then take more time to complete. Second, the progression of the execution is limited by the slowest cluster (gdx in this case). While this overhead remains lower than the number of sites, further experiments showed that it increases with the number of sites and is also greater for smaller problem classes. Indeed, a lower amount of computations leads to a greater impact of wide area communications. Finally the combination of process folding and scattering (SF-(2,2) to SF-(2,16)), the overhead are cumulated. If we divide the ratios to regular mode by the value obtained for S-2, we still observe that the execution time increases with the folding factor in a roughly linear way.

An interesting property of time-independent traces is exemplified by this set of experiments. A classical tracing tool such as TAU will produce traces full of erroneous timestamps in most of the scenarios. An off-line simulator using such traces will predict an execution time close to that of the corresponding acquisition scenario instead of the targeted **Regular** mode execution time. Preventing such a behavior would require to provide an accurate description of the acquisition platform along with the trace. However with time-independent traces, the simulated time is more or less the same whatever the acquisition scenario is. Slight variations lesser than 1% are observed that come from hardware counter accuracy issues.

6.3 Analysis of Trace Sizes

As mentioned in Section 2, the main challenge for off-line simulation is the large size of the traces. Here we propose to measure the size of the generated time-independent traces and compare it to the corresponding timed produced by TAU. Table 3 presents the obtained results for different instances of the LU benchmark.

		Trace size in MiB			
	#Processes	TAU	Time Independent	$\frac{\text{TAU}}{\text{SimGrid}}$	#Actions (in millions)
Class B	8	320.2	29.9	10.71	2.03
	16	716.5	72.6	9.87	4.87
	32	1,509	161.3	9.36	10.55
	64	3,166.1	344.9	9.18	22.73
Class C	8	508.2	48.4	10.5	3.23
	16	1,136.5	117	9.71	7.75
	32	2,393	256.8	9.32	16.79
	64	5,026.1	552.5	9.1	36.17

Table 3: Sizes of TAU and time-independent traces and number of actions for different instances of the LU benchmark.

We see that the time-independent traces are around 10 times smaller than the traces produced by TAU. The ratio slightly decreases as the number of processes increases. This shows the benefit of the event file of TAU that allows to factor the declaration of event names. The size of both types of traces grows linearly with the number of processes which is explained by the evolution of the number of traced events. We also see that the size of the time-independent traces grows from a constant factor of 1.6 from class B to class C which is also directly related to the number of actions.

6.4 Accuracy of Time-Independent Trace Replay

Figure 8: Comparison of simulated and actual execution time for the LU benchmark on the *bordereau* cluster.

Figure 8 shows the accuracy of the time-independent trace replay by comparing the actual execution time of the LU benchmark on the *bordereau* cluster for classes B and

C with the simulated time obtained with SIMGrid. We see that the trace replay is able to predict the correct trend of evolution of the execution time. However, the local relative error may be quite high (up to 51.5% for Class B on 64 processes) and not constant. This prevents our tool to provide predictions with a fixed interval of confidence.

These difficulties to match the experimental data are clearly related to the calibration of the simulator. More precisely it principally comes from the calibration of the flop rate that is not necessary with other off-line simulators that rely on timed traces. Indeed, the flop rate is not constant over the computation of a LU benchmark. Moreover, this flop rate does not even depend on the size of the computation. Improving the accuracy of the trace replay would imply to acquire more information on each computation during the calibration step to adapt the flop rate accordingly.

6.5 Acquiring a Large Trace

The analysis detailed in the previous sections were made on traces corresponding to small instances of the LU benchmark (up to 64 processes). Moreover this number of processes is smaller than the number of nodes in the compute clusters used for these experiments (respectively 93 and 186 nodes).

In order to demonstrate that the proposed approach can be used to assess the performance of MPI applications on clusters that are not available, we study the acquisition of traces for a large instance of the considered benchmark. We run our acquisition process on the *bordereau* cluster for a Class D instance executed on 1,024 processes. Such an instance is almost three times bigger than the number of cores ($93 \times 2 \times 2 = 372$) that this cluster comprises. To obtain this trace, we only used 32 nodes, i.e., 128 individual cores, and a folding factor of 8. This experiment then requires only about one third of the total amount of available resources.

It took less than 25 minutes to acquire (including extraction and gathering) the time-independent trace. Its size is 32.5 GiB, which is 7.8 times smaller than the TAU trace (252.5 GiB). When compressed with `gzip`, the time-independent trace takes 1.2 GiB.

6.6 Simulation Time

Figure 9 presents the evolution of the time needed to replay a time-independent trace as the number of processes increases for classes B and C instances of the LU benchmark. These timings were obtained on one node of the *bordereau* cluster.

We see that the replay time is directly related to the number of actions in the traces that is given in table 3. More precisely executing such a large number of actions (up to 36 millions for Class C on 64 processes) implies as much context switches between SIMGrid processes. Then the biggest part of this simulation time is spent in the system. For large traces such as that of the previous, the replay time increases to several hours. Solutions to reduce the number of context switches can be to write the simulator directly on top of the simulation, i.e., by bypassing the MSG API, or to distribute the simulated processes on several machines.

7 Conclusion and Future Work

In this paper, we proposed a new approach for the off-line simulation of MPI applications. Instead of relying on logs of execution that associate an event to a time-stamp,

Figure 9: Evolution of the trace replay time with the number of processes executing a LU benchmark.

we use time-independent traces as an input of our simulator. These traces contain only information about volumes of computation and communications. This allows for an acquisition process that is totally decoupled for the subsequent replay of the trace. Heterogeneous and distributed platforms can then be used to get an execution traces without impacting the quality of the simulation. Our simulator is built on top of the simulation kernel of the SIMGrid toolkit and thus benefits of its fast, scalable, and validated simulation kernel. We also rely on a well established tracing tool, TAU, as the first step of our acquisition process. In our experiments we have estimated the overhead required to produce a time-independent trace and showed the trace size reduction traces when compared to TAU. We also discussed the accuracy of the trace replay and the time needed to perform such simulations. The conclusion is that while some issues still have to be solved, decoupling acquisition and replay is sound and deserves further investigation.

As future work, we plan to solve the accuracy and simulation time issues for which we have hints of solutions. We also aim at exploring techniques to reduce the size of the traces, e.g., using a binary format. Finally we plan to compare off-line simulations results with those produced by on-line simulators.

Acknowledgments

This work is partially supported by the ANR project USS SimGrid (08-ANR-SEGI-022) and the CNRS PICS N° 5473. Experiments presented in this paper were carried out using the Grid'5000 experimental testbed, being developed under the INRIA AL-ADDIN development action with support from CNRS, RENATER and several Universities as well as other funding bodies (see <https://www.grid5000.fr>).

References

- [1] W. Gropp, E. Lusk, and A. Skjellum, *Using MPI: Portable Parallel Programming with the Message Passing Interface*, 2nd ed., ser. Scientific And Engineering Computation Series. MIT Press, 1999.
- [2] H. Casanova, A. Legrand, and M. Quinson, “SimGrid: a Generic Framework for Large-Scale Distributed Experiments,” in *Proc. of the 10th IEEE International Conference on Computer Modeling and Simulation*, Cambridge, UK, Mar. 2008.
- [3] P. Dickens, P. Heidelberger, and D. Nicol, “Parallelized Direct Execution Simulation of Message-Passing Parallel Programs,” *IEEE Trans. on Parallel and Distributed Systems*, vol. 7, no. 10, pp. 1090–1105, 1996.
- [4] R. Bagrodia, E. Deelman, and T. Phan, “Parallel Simulation of Large-Scale Parallel Applications,” *International Journal of High Performance Computing Applications*, vol. 15, no. 1, pp. 3–12, 2001.
- [5] R. Riesen, “A Hybrid MPI Simulator,” in *Proc. of the IEEE International Conference on Cluster Computing*, Barcelona, Spain, Sep. 2006.
- [6] G. Zheng, G. Kakulapati, and L. Kale, “BigSim: A Parallel Simulator for Performance Prediction of Extremely Large Parallel Machines,” in *Proc. of the 18th International Parallel and Distributed Processing Symposium*, Santa Fe, NM, Apr. 2004.
- [7] A. Snaveley, L. Carrington, N. Wolter, J. Labarta, R. Badia, and A. Purkayastha, “A Framework for Application Performance Modeling and Prediction,” in *Proc. of the ACM/IEEE Conference on Supercomputing (SC’02)*, Baltimore, MA, Nov. 2002.
- [8] E. León, R. Riesen, and A. Maccabe, “Instruction-Level Simulation of a Cluster at Scale,” in *Proc. of the International Conference for High Performance Computing and Communications*, Portland, OR, Nov. 2009.
- [9] S. Prakash, E. Deelman, and R. Bagrodia, “Asynchronous Parallel Simulation of Parallel Programs,” *IEEE Transactions on Software Engineering*, vol. 26, no. 5, pp. 385–400, 2000.
- [10] B. Penoff, A. Wagner, M. Tüxen, and I. Rüngeler, “MPI-NetSim: A network simulation module for MPI,” in *Proc. of the 15th International Conference on Parallel and Distributed Systems (ICPADS)*, Shenzhen, China, Dec. 2009.
- [11] T. Hoefler, C. Siebert, and A. Lumsdaine, “LogGOPSim - Simulating Large-Scale Applications in the LogGOPS Model,” in *Proc. of the ACM Workshop on Large-Scale System and Application Performance*, Chicago, IL, Jun. 2010, pp. 597–604.
- [12] M. Tikir, M. Laurenzano, L. Carrington, and A. Snaveley, “PSINS: An Open Source Event Tracer and Execution Simulator for MPI Applications,” in *Proc. of the 15th International EuroPar Conference*, ser. LNCS, vol. 5704, Delft, Aug. 2009, pp. 135–148.
- [13] A. Núñez, J. Fernández, J.-D. García, F. García, and J. Carretero, “New Techniques for Simulating High Performance MPI Applications on Large Storage Networks,” *Journal of Supercomputing*, vol. 51, no. 1, pp. 40–57, 2010.

- [14] J. Zhai, W. Chen, and W. Zheng, “PHANTOM: Predicting Performance of Parallel Applications on Large-Scale Parallel Machines Using a Single Node,” in *Proc. of the 15th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, Jan. 2010, pp. 305–314.
- [15] M.-A. Hermanns, M. Geimer, F. Wolf, and B. Wylie, “Verifying Causality between Distant Performance Phenomena in Large-Scale MPI Applications,” in *Proc. of the 17th Euromicro International Conference on Parallel, Distributed and Network-based Processing*, Weimar, Germany, Feb. 2009, pp. 78–84.
- [16] R. Badia, J. Labarta, J. Giménez, and F. Escalé, “Dimemas: Predicting MPI applications behavior in Grid environments,” in *Proc. of the Workshop on Grid Applications and Programming Tools*, 2003.
- [17] G. Zheng, T. Wilmarth, P. Jagadishprasad, and L. Kalé, “Simulation-Based Performance Prediction for Large Parallel Machines,” *Int. Journal of Parallel Programming*, vol. 33, no. 2-3, pp. 183–207, 2005.
- [18] S. Shende and A. Malony, “The Tau Parallel Performance System,” *International Journal of High Performance Computing Applications*, vol. 20, no. 2, pp. 287–311, 2006.
- [19] S. Browne, J. Dongarra, N. Garner, G. Ho, and P. Mucci, “A Portable Programming Interface for Performance Evaluation on Modern Processors,” *International Journal of High Performance Computing Applications*, vol. 14, no. 3, pp. 189–204, 2000.
- [20] Performance Research Lab, *TAU User Guide, chapter Tracing, TAU Trace Format Reader Library*, University of Oregon, <http://www.cs.uoregon.edu/research/tau/docs/newguide/ch06s02.html>.
- [21] P. Velho and A. Legrand, “Accuracy Study and Improvement of Network Simulation in the SimGrid Framework,” in *Proc. of the 2nd International Conference on Simulation Tools and Technique (SIMUTools'09)*, Roma, Italy, Mar. 2009.
- [22] M. Geimer, F. Wolf, B. Wylie, and B. Mohr, “A Scalable Tool Architecture for Diagnosing Wait States in Massively Parallel Applications,” *Parallel Computing*, vol. 35, no. 7, pp. 375–388, 2009.
- [23] R. Reussner, P. Sanders, and J. L. Träff, “SKaMPI: a Comprehensive Benchmark for Public Benchmarking of MPI,” *Scientific Programming*, vol. 10, no. 1, pp. 55–65, 2002.

Centre de recherche INRIA Nancy – Grand Est
LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399