

TOWARDS THE AUTOMATIC MAPPING OF ASSIST APPLICATIONS FOR THE GRID

Marco Aldinucci

*Computer Science Departement, University of Pisa
Largo Bruno Pontecorvo 3, I-56127 Pisa, Italy*

aldinuc@di.unipi.it

Anne Benoit

*LIP, Ecole Normale Supérieure de Lyon
46 Allée d'Italie, 69364 Lyon Cedex 07, France*

Anne.Benoit@ens-lyon.fr

Abstract One of the most promising technical innovations in present-day computing is the invention of grid technologies which harness the computational power of widely distributed collections of computers. However, the programming and optimisation burden of a low level approach to grid computing is clearly unacceptable for large scale, complex applications. The development of grid applications can be simplified by using high-level programming environments. In the present work, we address the problem of the mapping of a high-level grid application onto the computational resources. In order to optimise the mapping of the application, we propose to automatically generate performance models from the application using the process algebra PEPA. We target in this work applications written with the high-level environment ASSIST, since the use of such a structured environment allows us to automate the study of the application more effectively.

Keywords: high-level parallel programming, grid, ASSIST, PEPA, automatic model generation, skeletons.

1. Introduction

A grid system is a geographically distributed collection of possibly parallel, interconnected processing elements, which all run some form of common grid middleware (e.g. Globus services) [16]. The key idea behind grid-aware applications is to make use of the aggregate power of distributed resources, thus benefiting from a computing power that falls far beyond the current availability threshold in a single site. However, developing programs able to exploit this potential is highly programming intensive. Programmers must design concurrent programs that can execute on large-scale platforms that cannot be assumed to be homogeneous, secure, reliable or centrally managed. They must then implement these programs correctly and efficiently. As a result, in order to build efficient grid-aware applications, programmers have to address the classical problems of parallel computing as well as grid-specific ones:

1. *Programming*: code all the program details, take care about concurrency exploitation, among the others: concurrent activities set up, mapping/scheduling, communication/synchronisation handling and data allocation.

2. *Mapping & Deploying*: deploy application processes according to a suitable mapping onto grid platforms. These may be highly heterogeneous in architecture and performance. Moreover, they are organised in a cluster-of-clusters fashion, thus exhibiting different connectivity properties among all pairs of platforms.

3. *Dynamic environment*: manage resource unreliability and dynamic availability, network topology, latency and bandwidth unsteadiness.

Hence, the number and quality of problems to be resolved in order to draw a given QoS (in term of performance, robustness, etc.) from grid-aware applications is quite large. The lesson learnt from parallel computing suggests that any low-level approach to grid programming is likely to raise the programmer's burden to an unacceptable level for any real world application.

Therefore, we envision a layered, high-level programming model for the grid, which is currently pursued by several research initiatives and programming environments, such as ASSIST [22], eSkel [10], GrADS [20], ProActive [7], Ibis [21], Higher Order Components [13–14]. In such an environment, most of the grid specific efforts are moved from programmers to grid tools and run-time systems. Thus, the programmers have only the responsibility of organising the application specific code, while the programming tools (i.e. the compiling tools and/or the run-time systems) deal with the interaction with the grid, through collective protocols and services [15].

In such a scenario, the QoS and performance constraints of the application can either be specified at compile time or varying at run-time. In both cases, the run-time system should actively operate in order to fulfil QoS requirements of the application, since any static resource assignment may violate QoS constraints

due to the very uneven performance of grid resources over time. As an example, ASSIST applications exploit an autonomic (self-optimisation) behavior. They may be equipped with a QoS contract describing the degree of performance the application is required to provide. The ASSIST run-time environment tries to keep the QoS contract valid for the duration of the application run despite possible variations of platforms' performance at the level of grid fabric [6, 5]. The autonomic features of an ASSIST application rely heavily on run-time application monitoring, and thus they are not fully effective for application deployment since the application is not yet running. In order to deploy an application onto the grid, a suitable mapping of application processes onto grid platforms should be established, and this process is quite critical for application performance.

This problem can be addressed by defining a performance model of an ASSIST application in order to statically optimise the mapping of the application onto a heterogeneous environment, as shown in [1]. The model is generated from the source code of the application, before the initial mapping. It is expressed with the process algebra PEPA [18], designed for performance evaluation. The use of a stochastic model allows us to take into account aspects of uncertainty which are inherent to grid computing, and to use classical techniques of resolution based on Markov chains to obtain performance results. This static analysis of the application is complementary with the autonomic reconfiguration of ASSIST applications, which works on a dynamic basis. In this work we concentrated on the static part to optimise the mapping, while the dynamic management is done at run-time. It is thus an orthogonal but complementary approach.

Structure of the paper. The next section introduces the ASSIST high-level programming environment and its run-time support. Section 4.2 introduces the Performance Evaluation Process Algebra PEPA, which can be used to model ASSIST applications. These performance models help to optimise the mapping of the application. We present our approach in Section 4, and give an overview of future working directions. Finally, concluding remarks are given in Section 5.

2. The ASSIST environment and its run-time support

ASSIST (A Software System based on Integrated Skeleton Technology) is a programming environment aimed at the development of distributed high-performance applications [22, 3]. ASSIST applications should be compiled in binary packages that can be deployed and run on grids, including those exhibiting heterogeneous platforms. Deployment and run is provided through standard middleware services (e.g. Globus) enriched with the ASSIST run-time support.

2.1 The ASSIST coordination language

ASSIST applications are described by means of a coordination language, which can express arbitrary graphs of modules, interconnected by typed streams of data. Each stream realises a one-way asynchronous channel between two sets of endpoint modules: sources and sinks. Data items injected from sources are broadcast to all sinks. All data items injected into a stream should match the stream type.

Modules can be either sequential or parallel. A sequential module wraps a sequential function. A parallel module (*parmod*) can be used to describe the parallel execution of a number of sequential functions that are activated and run as *Virtual Processes* (VPs) on items arriving from input streams. The VPs may synchronise with the others through barriers. The sequential functions can be programmed by using a standard sequential language (C, C++, Fortran). A *parmod* may behave in a data-parallel (e.g. SPMD/for-all/apply-to-all) or task-parallel (e.g. farm) way and it may exploit a distributed shared state that survives the VPs lifespan. A module can nondeterministically accept from one or more input streams a number of input items according to a CSP specification included in the module [19]. Once accepted, each stream item may be decomposed in parts and used as function parameters to instantiate VPs according to the input and distribution rules specified in the *parmod*. The VPs may send items or parts of items onto the output streams, and these are gathered according to the output rules. Details on the ASSIST coordination language can be found in [22, 3].

2.2 The ASSIST run-time support

The ASSIST compiler translates a graph of modules into a network of processes. As sketched in Fig. 1, sequential modules are translated into sequential processes, while parallel modules are translated into a parametric (w.r.t. the parallelism degree) network of processes: one *Input Section Manager* (ISM), one *Output Section Manager* (OSM), and a set of *Virtual Processes Managers* (VPMs, each of them running a set of Virtual Processes). The ISM implements a CSP interpreter that can send data items to VPMs via collective communications. The number of VPMs gives the actual parallelism degree of a *parmod* instance. Also, a number of processes are devoted to application dynamic QoS control, e.g. a *Module Adaptation Manager* (MAM), and an *Application Manager* (AM) [6, 5].

The processes that compose an ASSIST application communicate via ASSIST support channels. These can be implemented on top of a number of grid middleware communication mechanisms (e.g. shared memory, TCP/IP, Globus, CORBA-IIOP, SOAP-WS). The suitable communication mechanism between each pair of processes is selected at launch time depending on the mapping of the processes.

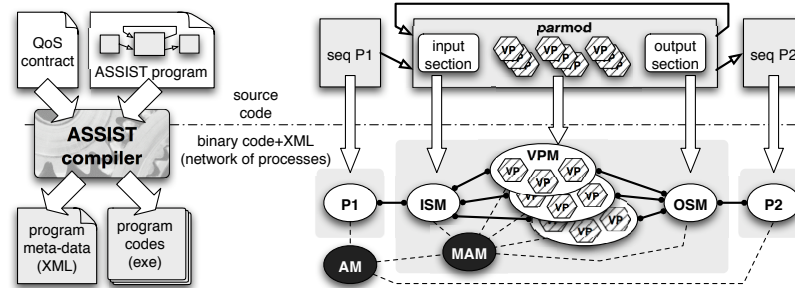


Figure 1. An ASSIST application and a QoS contract are compiled in a set of executable codes and its meta-data [3]. This information is used to set up a processes network at launch time.

2.3 Towards fully grid-aware applications

ASSIST applications can already cope with platform heterogeneity [2], either in space (various architectures) or in time (varying load) [6]. These are definite features of a grid, however they are not the only ones. Grids are usually organised in sites on which processing elements are organised in networks with private addresses allowing only outbound connections. Also, they are often fed through job schedulers. In these cases, setting up a multi-site parallel application onto the grid is a challenge in its own right (irrespective of its performance). Advance reservation, co-allocation, multi-site launching are currently hot topics of research for a large part of the grid community. Nevertheless, many of these problems should be targeted at the middleware layer level and they are largely independent of the logical mapping of application processes on a suitable set of resources, given that the mapping is consistent with deployment constraints.

In our work, we assume that the middleware level supplies (or will supply) suitable services for co-allocation, staging and execution. These are actually the minimal requirements in order to imagine the bare existence of any non-trivial, multi-site parallel application. Thus we can analyse how to map an ASSIST application, assuming that we can exploit middleware tools to deploy and launch applications [12].

3. Introduction to performance evaluation and PEPA

In this section, we briefly introduce the Performance Evaluation Process Algebra PEPA [18], with which we can model an ASSIST application. The use of a process algebra allows us to include the aspects of uncertainty relative to both the grid and the application, and to use standard methods to easily and quickly obtain performance results.

The PEPA language provides a small set of combinators. These allow language terms to be constructed defining the behavior of components, via the activities they undertake and the interactions between them. We can for instance define constants, express the sequential behavior of a given component, a choice between different behaviors, and the direct interaction between components. Timing information is associated with each activity. Thus, when enabled, an activity $a = (\alpha, r)$ will delay for a period sampled from the negative exponential distribution which has parameter r . If several activities are enabled concurrently, either in competition or independently, we assume that a *race condition* exists between them.

The dynamic behavior of a PEPA model is represented by the evolution of its components, as governed by the operational semantics of PEPA terms [18]. Thus, as in classical process algebra, the semantics of each term is given via a labelled *multi-transition* system (the multiplicity of arcs are significant). In the transition system a state corresponds to each syntactic term of the language, or *derivative*, and an arc represents the activity which causes one derivative to evolve into another. The complete set of reachable states is termed the *derivative set* and these form the nodes of the *derivation graph*, which is formed by applying the semantic rules exhaustively. The derivation graph is the basis of the underlying Continuous Time Markov Chain (CTMC) which is used to derive performance measures from a PEPA model. The graph is systematically reduced to a form where it can be treated as the state transition diagram of the underlying CTMC. Each derivative is then a state in the CTMC. The *transition rate* between two derivatives P and Q in the derivation graph is the rate at which the system changes from behaving as component P to behaving as Q . Examples of derivation graphs can be found in [18].

It is important to note that in our models the rates are represented as random variables, not constant values. These random variables are exponentially distributed. Repeated samples from the distribution will follow the distribution and conform to the mean but individual samples may potentially take any positive value. The use of such distribution is quite realistic and it allows us to use standard methods on CTMCs to readily obtain performance results. There are indeed several methods and tools available for analysing PEPA models. Thus, the PEPA Workbench [17] allows us to generate the state space of a PEPA model and the infinitesimal generator matrix of the underlying Markov chain. The state space of the model is represented as a sparse matrix. The PEPA Workbench can then compute the steady-state probability distribution of the system, and performance measures such as throughput and utilisation can be directly computed from this.




Figure 2. Graph representation of our example application.

4. Performance models of ASSIST applications

PEPA can easily be used to model an ASSIST application since such applications are based on stream communications, and the graph structure deduced from these streams can be modelled with PEPA. Given the probabilistic information about the performance of each of the ASSIST modules and streams, we then aim to find information about the global behavior of the application, which is expressed by the steady-state of the system. The model thus allows us to predict the run-time behavior of the application in the long time run, taking into account information obtained from a static analysis of the program. This behavior is not known in advance, it is a result of the PEPA model.

4.1 The ASSIST application

As we have seen in Section 2, an ASSIST application consists of a series of modules and streams connecting the modules. The structure of the application is represented by a graph, where the modules are the nodes and the streams the arcs.

We illustrate in this paper our modeling process on an example of a graph, but the process can be easily generalized to any ASSIST applications since the information about the graph can be extracted directly from ASSIST source code, and the model can be generated automatically from the graph.

A model of a data mining classification algorithm has been presented in [1], as well as the corresponding ASSIST source code. For the purpose of our methodology and in order to generalize our approach, we concentrate here only on the graph of an application.

The graph of the application that we consider in this paper is similar to the one of [1], consisting of four modules. Figure 2 represents the graph of this application.

4.2 The PEPA model

Each ASSIST module is represented as a PEPA component, and the different components are synchronised through the streams of data to model the overall

application. The performance results obtained are the probabilities to be in either of the states of the system. From this information, we can determine the bottleneck of the system and decide the best way to map the application onto the available resources.

The PEPA model is generated automatically from the ASSIST source code, during a pre-compilation phase. The information required for the generation is provided by the user directly in the source code, and particularly the rates associated to the different activities of the PEPA model. These rates are related to the theoretical complexity of the modules and of the communications. In particular, rates of the communications depend on: a) the speed of the links and b) data size and communications frequencies. A module may include a parallel computation, thus its rate depend on a) computing power of the platforms running the module and b) parallel computation complexity, its size, its parallel degree, and its speedup.

Observe that aspect a) of both modules and communications rates strictly depend on mapping, while aspect b) is much more dependent by application logical structure and algorithms.

The PEPA components of the modules are shown in Fig. 3. The modules are working in a sequential way: the module MX ($X = 1..4$) is initially in the state $MX1$, waiting for data on its input streams. Then, in the state $MX2$, it processes the piece of data and evolves to its third state $MX3$. Finally, the module sends the output data on its output streams and goes back into its first state.

The system evolves from one state to another when an activity occurs. The activity sX ($X = 1..4$) represents the transfer of data through the stream X , with the associated rate λ_X . The rate reflects the complexity of the communication. The activity pX ($X = 1..4$) represents the processing of a data by module MX , which is done at a rate μ_X . These rates are related to the theoretical complexity of the modules.

The overall PEPA model is then obtained by a collaboration of the different modules in their initial states: $M11 \bowtie_{s_1} M21 \bowtie_{s_2, s_3} M31 \bowtie_{s_4} M41$.

The performance results obtained are the probability to be in either of the states of the system. We compute the probability to be waiting for a processing activity pX , or to wait for a transfer activity sX . From this information, we can determine the bottleneck of the system and decide the best way to map the application onto the available resources.

4.3 Automatic generation of the model

To allow an automatic generation of the PEPA model from the ASSIST source code, we ask the user to provide some information directly in the main procedure of the application. This information must specify the rates of the different activities of the PEPA model. We are interested in the relative computational

$$\begin{aligned}
M11 &\stackrel{\text{def}}{=} M12 \\
M12 &\stackrel{\text{def}}{=} (p1, \mu_1).M13 \\
M13 &\stackrel{\text{def}}{=} (s1, \lambda_1).M11 \\
\\
M21 &\stackrel{\text{def}}{=} (s1, \top).M22 + (s2, \top).M22 \\
M22 &\stackrel{\text{def}}{=} (p2, \mu_2).M23 \\
M23 &\stackrel{\text{def}}{=} (s3, \lambda_3).M21 + (s4, \lambda_4).M21 \\
\\
M31 &\stackrel{\text{def}}{=} (s3, \top).M32 \\
M32 &\stackrel{\text{def}}{=} (p3, \mu_3).M33 \\
M33 &\stackrel{\text{def}}{=} (s2, \lambda_2).M31 \\
\\
M41 &\stackrel{\text{def}}{=} (s4, \top).M42 \\
M42 &\stackrel{\text{def}}{=} (p4, \mu_4).M43 \\
M43 &\stackrel{\text{def}}{=} M41
\end{aligned}$$

Figure 3. PEPA model for the example

and communication costs of the different parts of the system, but we define numerical values to allow a numerical resolution of the PEPA model.

The complexity of the modules depends on the number of computations done, and also on the degree of parallelism used for a parallel module. It is directly related to the time needed to compute one input. The rates associated with the streams depends on the amount of the data transiting on each stream. In ASSIST, the object transiting on the stream is often a reference to the real object, since the actual data is available in a shared memory, and this is beyond the scope of our PEPA model.

This information is defined directly in the ASSIST source code of the application, by calling a `rate` function, which takes as a parameter the name of the modules and streams. This function should be called once for each module and each stream to fix the rates of the corresponding PEPA activities.

The PEPA model is generated during a precompilation of the source code of ASSIST. The parser identifies the `main` procedure and extracts the useful information from it: the modules and streams, the connections between them, and the rates of the different activities. The main difficulty consists in identifying the schemes of input and output behavior in the case of several streams. This information can be found in the input and output section of the `parmod` code. Regarding the input section, the parser looks at the guards. Details on the different types of guards can be found in [22, 3].

As an example, a disjoint guard means that the module takes input from either of the streams when some data arrives. This is translated by a choice in the PEPA model, as illustrated in our example. However, some more complex behavior may also be expressed, for instance the parmod can be instructed to start executing only when it has data from both streams. In this case, the PEPA model is changed with some sequential composition to express this behavior. For example, $M21 \stackrel{\text{def}}{=} (s1, \top).(s2, \top).M22 + (s2, \top).(s1, \top).M22$.

Another problem may arise from the variables in guards, since these may change the frequency of accessing data from a stream. Since the variables may depend on the input data, we cannot automatically extract static information from them. They are currently ignored, but we plan to address this problem by asking the programmer to provide the relative frequency of the guard. The considerations for the output section are similar.

4.4 Performance results

Once the PEPA model has been generated, performance results can be obtained easily with the PEPA Workbench [17]. Some additional information is generated in the PEPA source code to specify the performance results that we are interested in. This information is the following:

```

moduleM1 = 100 * {M12 || ** || ** || ** };
moduleM2 = 100 * {** || M22 || ** || ** };
moduleM3 = 100 * {** || ** || M32 || ** };
moduleM4 = 100 * {** || ** || ** || M42};

stream1 = 100 * {M13 || M21 || ** || ** };
stream2 = 100 * {** || M21 || M33 || ** };
stream3 = 100 * {** || M23 || M31 || ** };
stream4 = 100 * {** || M23 || ** || M41};

```

The expression in brackets describes the states of the PEPA model corresponding to a particular state of the system. For each module MX ($X = 1..4$), the result moduleMX corresponds to the percentage of time spent waiting to process and processing this module. The steady-state probability is multiplied by 100 for readability and interpretation reasons. A similar result is obtained for each stream.

We expect the complexity of the PEPA model to be quite simple and the resolution straightforward for most of the ASSIST applications. In our example, the PEPA model consists in 36 states and 80 transitions, and it requires less than 0.1 seconds to generate the state space of the model and to compute the steady state solution, using the linear biconjugate gradient method [17].

Experiment 1 For the purpose of our example, we choose the following rates, meaning that the module *M3* is computationally more intensive than the other modules. In our case, *M3* has an average duration of 1 sec. compared to 0.01 sec. for the others: $\mu_1 = 100; \mu_2 = 100; \mu_3 = 1; \mu_4 = 100;$

The rates for the streams correspond to an average duration of 0.1 sec: $\lambda_1 = 10; \lambda_2 = 10; \lambda_3 = 10; \lambda_4 = 10;$

The results for this example are shown in Table 1 (row *Case 1*).

These results confirm the fact that most of the time is spent in module *M3*, which is the most computationally demanding. Moreover, module *M1* (respectively *M4*) spends most of its time waiting to send data on *s1* (respectively waiting to receive data from *s4*). *M2* is computing quickly, and this module is often receiving/sending from stream *s2/s3* (little time spent waiting on these streams in comparison with streams *s1/s4*).

If we study the computational rate, we can thus decide to map *M3* alone on a powerful processor because it has the highest value between the different steady states probabilities of the modules. One should be careful to map the streams *s1* and *s4* onto sufficiently fast network links to increase the overall throughput of the network. A mapping that performs well can thus be deduced from this information, by adjusting the reasoning to the architecture of the available system.

Experiment 2 We can reproduce the same experiment but for a different application: one in which there are a lot of data to be transferred inside the loop. Here, for one input on *s1*, the module *M2* makes several calls to the server *M3* for computations. In this case, the rates of the streams are different, for instance $\lambda_1 = \lambda_4 = 1000$ and $\lambda_2 = \lambda_3 = 1$.

The results for this experiment are shown in Table 1 (row *Case 2*). In this table, we can see that *M3* is quite idle, waiting to receive data 89.4% of the time (i.e. this is the time it is not processing). Moreover, we can see in the stream results that *s2* and *s3* are busier than the other streams. In this case a good solution might be to map *M2* and *M3* on to the same cluster, since *M3* is no longer the computational bottleneck. We could thus have fast communication links for *s2* and *s3*, which are demanding a lot of network resources.

Table 1. Performance results for the example.

	Modules				Streams			
	M1	M2	M3	M4	s1	s2	s3	s4
Case 1	4.2	5.1	67.0	4.2	47.0	6.7	6.7	47.0
Case 2	52.1	52.2	10.6	52.1	5.2	10.6	10.6	5.2

4.5 Analysis summary

As mentioned in Section 4.2, PEPA rates model both aspects strictly related to the mapping and to the application's logical structure (such as algorithms implemented in the modules, communication patterns and size). The predictive analysis conducted in this work provides performance results which are related only to the application's logical behavior. On the PEPA model this translates on the assumption that all sites includes platforms with the same computing power, and all links have an uniform speed. In other words, we assume to deal with a homogeneous grid to obtain the relative requirements of power among links and platforms. This information is used as a hint for the mapping in a heterogeneous grid.

It is of value to have a general idea of a good mapping solution for the application, and this reasoning can be easily refined with new models including the mapping peculiarities, as demonstrated in our previous work [1]. However, the modeling technique exposed in the present paper allows us to highlight individual resources (links and processors) requirements, that are used to label the application graph.

These labels represent the expected relative requirements of each module (stream) with respect to other modules (streams) during the application run. In the case of a module the described requirement can be interpreted as the aggregate power of the site on which it will be mapped. On the other hand, a stream requirement can be interpreted as the bandwidth of the network link on which it will be mapped. The relative requirements of parmods and streams may be used to implement mapping heuristics which assign more demanding parmods to more powerful sites, and more demanding streams to links exhibiting higher bandwidths. When a fully automatic application mapping is not required, modules and streams requirements can be used to drive a user-assisted mapping process.

Moreover, each parmod exhibits a structured parallelism pattern (a.k.a. skeleton). In many cases, it is thus possible to draw a reliable relationship between the site fabric level information (number and kind of processors, processors and network benchmarks) and the expected aggregate power of the site running a given parmod exhibiting a parallelism pattern [5, 4, 9]. This may enable the development of a mapping heuristic, which needs only information about sites fabric level information, and can automatically derive the performance of a given parmod on a given site.

The use of models taking into account both of the system architecture characteristics can then eventually validate this heuristic, and give expected results about the performance of the application for a specified mapping.

4.6 Future work

The approach described here considers the ASSIST modules as blocks and does not model the internal behavior of each module. A more sophisticated approach might be to consider using known models of individual modules and to integrate these with the global ASSIST model, thus providing a more accurate indication of the performance of the application. At this level of detail, distributed shared memory and external services (e.g. DB, storage services, etc) interactions can be taken into account and integrated to enrich the network of processes with dummy nodes representing external services. PEPA models have already been developed for pipeline or deal skeletons [8–9], and we could integrate such models when the parmod module has been adapted to follow such a pattern.

Analysis precision can be improved by taking into account historical (past runs) or synthetic (benchmark) performance data of individual modules and their communications. This kind of information should be scaled with respect to the expected performances of fabric resources (platform and network performances), which can be retrieved via the middleware information system (e.g. Globus GIS).

We believe that this approach is particularly suitable for modeling applications that can be described by a graph, not just ASSIST applications (such as applications described in the forthcoming CoreGrid Grid Component Model [11]). In particular the technique described here helps to derive some information about the pressure (on modules and links) within a loop of the graph. Loops are quite common patterns; they can be used to describe simple interactions between modules (e.g. client-server RPC behavior) or mutual recursive dependency between modules. These two cases lead to very different behaviors in term of pressure or resources within the loop; in the former case this pressure is variable over time.

The mapping decision is inherently a static process, and especially for loops in the graph, it is important to make decisions on the expected common case. This is modeled by the PEPA steady state probabilities, that indeed try to give some static information on dynamic processes. Observe that PEPA is known to give much more precise information compared to other known methods, such as networks of queues, which cannot model finite buffering in queues, but it is possible with PEPA. Clearly this is important, particularly for loops within the graph.

5. Conclusions

In this paper we have presented a method to automatically generate PEPA models from an ASSIST application with the aim of improving the mapping of the application. This is an important problem in grid application optimisation.

It is our belief that having an automated procedure to generate PEPA models and obtain performance information may significantly assist in taking mapping decisions. However, the impact of this mapping on the performance of the application with real code requires further experimental verification. This work is ongoing, and is coupled with further studies on more complex applications.

This ongoing research is a collaboration between two CoreGRID partners: the University of Pisa, Italy (WP3 - Programming Model), and the ENS (CNRS) in Lyon, France (WP6 - Institute on Resource Management and Scheduling).

Acknowledgments

This work has been partially supported by Italian national FIRB project no. RBNE01KNFP *GRID.it*, by Italian national strategic projects *legge 449/97* No. 02.00470.ST97 and 02.00640.ST97, and by the FP6 Network of Excellence CoreGRID funded by the European Commission (Contract IST-2002-004265).

References

- [1] M. Aldinucci and A. Benoit. Automatic mapping of ASSIST applications using process algebra. Technical report TR-0016, CoreGRID, Oct. 2005.
- [2] M. Aldinucci, S. Campa, M. Coppola, S. Magini, P. Pesciullesi, L. Potiti, R. Ravazzolo, M. Torquati, and C. Zoccolo. Targeting heterogeneous architectures in ASSIST: Experimental results. In M. Danelutto, M. Vanneschi, and D. Laforenza, editors, *Proc. of 10th Intl. Euro-Par 2004 Parallel Processing*, volume 3149 of *LNCS*, pages 638–643. Springer Verlag, Aug. 2004.
- [3] M. Aldinucci, M. Coppola, M. Danelutto, M. Vanneschi, and C. Zoccolo. ASSIST as a research framework for high-performance grid programming environments. In J. C. Cunha and O. F. Rana, editors, *Grid Computing: Software environments and Tools*, chapter 10, pages 230–256. Springer Verlag, Jan. 2006.
- [4] M. Aldinucci, M. Danelutto, J. D nnweber, and S. Gorlatch. Optimization techniques for skeletons on grid. In L. Grandinetti, editor, *Grid Computing and New Frontiers of High Performance Processing*, volume 14 of *Advances in Parallel Computing*, chapter 2, pages 255–273. Elsevier, Oct. 2005.
- [5] M. Aldinucci, M. Danelutto, and M. Vanneschi. Autonomic QoS in ASSIST grid-aware components. In *Proc. of Intl. Euromicro PDP 2006: Parallel Distributed and network-based Processing*, pages 221–230, Montb liard, France, Feb. 2006. IEEE.
- [6] M. Aldinucci, A. Petrocelli, E. Pistoletti, M. Torquati, M. Vanneschi, L. Veraldi, and C. Zoccolo. Dynamic reconfiguration of grid-aware applications in ASSIST. In J. C. Cunha and P. D. Medeiros, editors, *Proc. of 11th Intl. Euro-Par 2005 Parallel Processing*, volume 3648 of *LNCS*, pages 771–781. Springer Verlag, Aug. 2005.
- [7] F. Baude, D. Caromel, and M. Morel. On hierarchical, parallel and distributed components for grid programming. In V. Getov and T. Kielmann, editors, *Proc. of the Intl. Workshop on Component Models and Systems for Grid Applications*, CoreGRID series, pages 97–108, Saint-Malo, France, Jan. 2005. Springer Verlag.
- [8] A. Benoit, M. Cole, S. Gilmore, and J. Hillston. Evaluating the performance of skeleton-based high level parallel programs. In M. Bubak, D. van Albada, P. Sloot, and J. Dongarra,

- editors, *The Intl. Conference on Computational Science (ICCS 2004), Part III*, LNCS, pages 299–306. Springer Verlag, 2004.
- [9] A. Benoit, M. Cole, S. Gilmore, and J. Hillston. Scheduling skeleton-based grid applications using PEPA and NWS. *The Computer Journal*, 48(3):369–378, 2005.
 - [10] M. Cole. Bringing Skeletons out of the Closet: A Pragmatic Manifesto for Skeletal Parallel Programming. *Parallel Computing*, 30(3):389–406, 2004.
 - [11] CoreGRID NoE deliverable series, Institute on Programming Model. *Deliverable D.PM.02 – Proposals for a Grid Component Model*, Nov. 2005.
 - [12] M. Danelutto, M. Vanneschi, C. Zoccolo, N. Tonellotto, S. Orlando, R. Baraglia, T. Fagni, D. Laforenza, and A. Paccosi. HPC application execution on grids. In V. Getov, D. Laforenza, and A. Reinefeld, editors, *Future Generation Grids*, CoreGRID series, pages 263–282. Springer Verlag, Nov. 2005.
 - [13] J. Dünweber and S. Gorch. HOC-SA: A grid service architecture for higher-order components. In *IEEE Intl. Conference on Services Computing, Shanghai, China*, pages 288–294. IEEE Computer Society Press, Sept. 2004.
 - [14] J. Dünweber, S. Gorch, S. Campa, M. Aldinucci, and M. Danelutto. Behavior customization of parallel components application programming. Technical Report TR-0002, Institute on Programming Model, CoreGRID - Network of Excellence, Apr. 2005.
 - [15] I. Foster, C. Kesselman, and S. Tuecke. The anatomy of the Grid: Enabling scalable virtual organization. *The Intl. Journal of High Performance Computing Applications*, 15(3):200–222, Fall 2001.
 - [16] I. Foster and C. Kesselmann, editors. *The Grid 2: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann, Dec. 2003.
 - [17] S. Gilmore and J. Hillston. The PEPA Workbench: A Tool to Support a Process Algebra-based Approach to Performance Modelling. In *Proc. of the 7th Int. Conf. on Modelling Techniques and Tools for Computer Performance Evaluation*, number 794 in LNCS, pages 353–368, Vienna, May 1994. Springer-Verlag.
 - [18] J. Hillston. *A Compositional Approach to Performance Modelling*. Cambridge University Press, 1996.
 - [19] C. A. R. Hoare. Communicating Sequential Processes. *Communications of ACM*, 21(8):666–677, Aug. 1978.
 - [20] S. Vadhiyar and J. Dongarra. Self adaptability in grid computing. *Concurrency & Computation: Practice & Experience*, 17(2–4):235–257, 2005.
 - [21] R. V. van Nieuwpoort, J. Maassen, G. Wrzesinska, R. Hofman, C. Jacobs, T. Kielmann, and H. E. Bal. Ibis: a flexible and efficient Java-based grid programming environment. *Concurrency & Computation: Practice & Experience*, 17(7–8):1079–1107, 2005.
 - [22] M. Vanneschi. The programming model of ASSIST, an environment for parallel and distributed portable applications. *Parallel Computing*, 28(12):1709–1732, Dec. 2002.