

# Adaptive Computation over Dynamic and Heterogeneous Networks

Kaoutar El Maghraoui<sup>1</sup>, Joseph E. Flaherty<sup>1</sup>, Boleslaw K. Szymanski<sup>1</sup>, James D. Teresco<sup>2</sup>, and Carlos Varela<sup>1</sup>

<sup>1</sup> Rensselaer Polytechnic Institute, Troy, NY 12180, USA,  
szymansk@cs.rpi.edu,  
<http://www.cs.rpi.edu/>

<sup>2</sup> Williams College, Williamstown, MA 01267, USA

**Abstract.** Over the last two decades, efficient message passing libraries have been developed for parallel scientific computation. Concurrently, programming languages have been created supporting dynamically reconfigurable distributed systems over the heterogeneous Internet. In this paper, we introduce SALSAMPI, an actor programming language approach to scientific computing that extends current MPI implementation through a checkpointing and migration API, a SALSAMPI device, and a runtime system that manages both periodic checkpoints and process or application migration. The goal is to enable dynamic network reconfiguration and load balancing without sacrificing application performance or requiring extensive code modifications. As driving technology for this effort of unifying parallel and distributed computing, we plan to use adaptive solvers of partial differential equations. Fields as diverse as fluid dynamics, material science, biomechanics, and ecology make use of parallel adaptive computation, but target architectures have traditionally been supercomputers and tightly-coupled clusters. SALSAMPI is intended to allow these computations to make efficient use of more distributed and dynamic computing resources.

## 1 Introduction

Large-scale scientific and engineering applications involving the solution of partial differential equations are among the most demanding computational problems, arising in fields including fluid dynamics, material science, biomechanics, and ecology. Adaptivity, where meshes and methods are automatically adjusted to achieve specified levels of solution accuracy, and parallelism are essential tools to solve modern multi-dimensional transient problems. The usual approach to these problems is to distribute a discretization (mesh) of the domain across cooperating processors, then to compute a solution, appraising its accuracy using error estimates at each step. If the solution is accepted, the computation proceeds to the next step. Otherwise, the discretization is refined adaptively, and work is redistributed, if necessary, to correct for any load imbalance introduced by the adaptive step. The adaptive strategies automatically refine, coarsen, and/or relocate meshes and may also change the method with a goal of obtaining a solution to a prescribed level of accuracy as quickly as possible [1]. Adaptivity makes automatic (compiler-directed) parallelization difficult, so dynamic partitioning and load balancing

procedures become necessary since the locations where meshes must be refined or simpler methods replaced by more complex ones are not known *a priori* and are determined as part of the solution process.

The importance of the applications and, perhaps, the cost of access to supercomputers have led to proliferation of solution strategies on other architectures including PC clusters and, most recently, grids [2, 3]. Target architectures range from small clusters to the largest supercomputers with interprocessor communication ranging from shared memory to wide-area networks.

As discussed in citeours in this volume, SALSA open source actor programming language and IO middleware provide distribution *transparency* to scientific programmers and support *efficient* efficient message passing. The main contribution of this paper is the SALSA-MPI middleware that supports dynamic partitioning and load balancing for parallel adaptive partial differential equation solvers and computation fault-tolerance via data and process migration and replication. When fully developed, SALSA-MPI will provide a fully integrated software framework linking the applications layer (programmer interface) with the middleware layer, so that adaptivity and transparency can be simultaneously and efficiently achieved.

## 2 Parallel Adaptive Scientific Computation

The adaptive software base developed at Rensselaer's Scientific Computation Research Center executes in serial and parallel computational environments [4, 5]. An adaptive approach with explicit domain decomposition has been used successfully by many software packages for classical finite element [6], finite volume [7], and discontinuous Galerkin (DGM) [8, 9] methods. *DG* [9] is a software package that implements a parallel adaptive DGM using the *Algorithm Oriented Mesh Database* (AOMD) [10] mesh structures and services. AOMD supports a variety of mesh representations, including hybrid meshes. It is written in C++ using the Standard Template Library [11] for computations and the Message Passing Interface (MPI) [12] for communication. *DG* is used to solve a wide range of problems including Rayleigh-Taylor flow instabilities [9]. Distributed AOMD meshes [5] use the *Rensselaer Partition Model* (RPM) [13] to aid in data distribution and migration.

## 3 Programming Abstractions and Technology for Dynamic Grids

The Java [14] platform – which includes the definition of a programming language, a virtual machine, and a set of libraries providing high-level application programming interfaces (API) – is a step forward in portable distributed software engineering. In particular, Java's support for concurrent and distributed programming includes multithreading and remote method invocation APIs. While Java's main perceived drawback is its lack of performance, due to its bytecode interpretation overhead, recent advances in JIT (Just In Time) compilation and adaptive compilation make Java a very attractive platform for scientific applications [15].

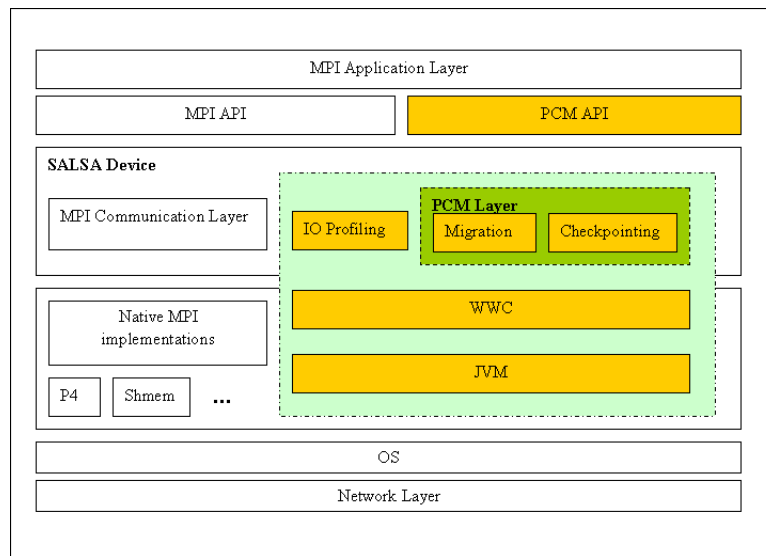
Unfortunately, Java's object model inhibits dynamic system reconfigurability for three reasons: passive objects, shared memory, and non-universal naming [16]. First, an

object can have multiple concurrent active threads within its body, drastically limiting its ability to migrate. Second, objects can share memory, requiring expensive cache coherence protocols upon migration. Last, objects do not have “universal” identifiers that enable references to persist upon object migration.

SALSA [17] is an actor-oriented programming language with high-level constructs for remote messaging, universal naming, migration, and coordination. SALSA programs are compiled into Java code, allowing a heterogeneous network of physical machines to be viewed as a homogeneous network of Java virtual machines. The WWC (World-Wide Computer) run-time architecture consists of naming servers and virtual machines running as Java applications on different Internet nodes. The virtual machines, called *theaters*, provide an environment for execution of universal actors using local resources. High-level programming language abstractions enable actors to create remote communication links with peer actors running on other WWC theaters. Furthermore, actors can easily migrate with their full state to other WWC theaters as they become available, supporting dynamic load balancing and scalability. The naming servers keep track of universal actor locators, so that communication remains transparent to actor location and migration.

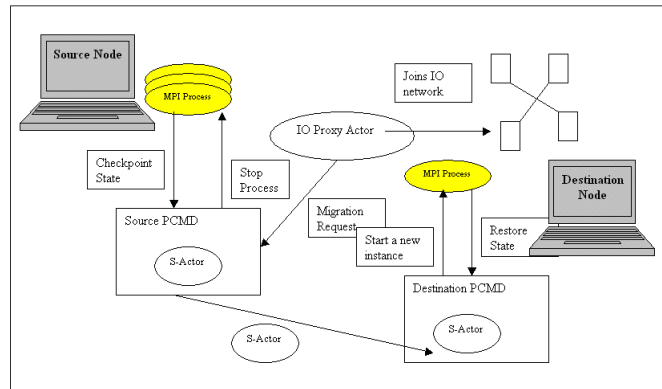
## 4 SALSA-MPI

### 4.1 SALSA-MPI Architecture



**Fig. 1.** SALSA-MPI Architecture.

The SALSA/IO architecture [18] consists of an actor-oriented programming language (SALSA), a distributed runtime environment (WWC), and a middleware infrastructure for autonomous load balancing (IO: Internet Operating System). This infrastructure is highly dynamic through its support of 1) resource profiling, 2) migration of actors to optimal execution environments, and 3) dynamic addition and removal of nodes in the computation. SALSA-MPI provides an extension to the MPI standard that allows MPI processes to run on dynamic grid environments. This is achieved through the use of the IO middleware that supports dynamic reconfiguration and load balancing. Figure 2 shows the proposed SALSA-MPI architecture. The SALSA device runs on top of vendor supplied MPI implementations. It uses intra-machine vendor supplied MPI implementations and inter-machine TCP communication. The SALSA device provides also an interface to SALSA/IO proxy actors, which act as profiling actors in the IO network. Every SALSA-MPI proxy-actor has a profiling actor and a decision agent. The profiling actor keeps monitoring the running MPI processes communication, the participating nodes' memory, CPU, and network resources. This monitoring is fed periodically to the decision agent. As more resources join or leave the computation, the decision agent tries to reconfigure the running MPI application by migrating one or more MPI processes in order to improve performance. This architecture supports intra-cluster process migration and inter-cluster application migration. It is expensive to spread MPI processes across clusters since they are usually very tightly coupled and they communicate a lot between themselves. Therefore, it is important to collocate all the running MPI process in a given application within one cluster. Inter-cluster Application migration on the other hand, could improve significantly the performance if the current cluster experiences failures or increased load.



**Fig. 2.** The different interactions between the components of the SALSA-MPI framework.

To support migration, we propose an application-level checkpointing API called PCM (Process Checkpointing and Migration) and a runtime system called PCMD (Pro-

cess Checkpointing and Migration Daemon). Few PCM calls need to be inserted in MPI programs to specify the data that need to be checkpointed, and to restore the process to its previous state in case of migration. This library is semi-transparent because the user does not have to worry about when or how checkpointing and restoration is done. The underlying PCMD infrastructure takes care of all the checkpointing and migration details.

## 4.2 The PCM API

The PCM API consists of a set of function calls that allow MPI programs to be dynamically reconfigurable. The PCM can be used with any iterative MPI application. The PCM library consists of set of routines with the following functionalities:

- Periodic checkpoints of MPI processes or MPI application data.
- Storage of the checkpointed data in a PCMD daemon (either locally or remotely).
- Restoration of a previously checkpointed data.
- Suspension, restart, or migration of an MPI process or an MPI application.
- Periodic probing of the status of an MPI application or an MPI process

Vadhiyar et al. have devised a similar approach through their SRS library [19]. Our approach is different in that our architecture proposes both process and application migration. In addition we use the SALSA-IO middleware to trigger reconfigurability and load balancing when necessary.


## 4.3 PCM Daemons

PCMD daemons need to be launched in every node that might join the MPI parallel computation. The PCMD has a port number on which it listens to incoming requests. It interacts with the running MPI processes, with the IO proxy actor, and with remote PCMD daemons. When the IO decision agent chooses a process for migration, it sends a migration request to the PCMD, which creates a shadow actor (S-Actor) for the migrating process, redirects all messages sent to the migrating process (MP). The MP initiates checkpoints at the PCMD daemon. The S-Actor serializes the state of the MP and migrates to the new destination. On the remote PCMD, the S-Actor starts a new instance of the MP, stores the states of the MP, notifies the source node to stop message redirections and terminates. Once the MP is restarted, it contacts the local PCMD to restore its state. Figure 2 illustrates the interactions between the PCM Daemons and the other components of the SALSA-MPI framework.

## 4.4 Preliminary Results

The experimental testbed consisted of two clusters at Rensselaer Polytechnic Institute: cluster A (the CS cluster) consisting of 20 256MB memory SUN 10 machines and cluster B (the Netfinity cluster) consisting of 40 128MB memory 900Mhz singel processor IBM Netifnity machines. The used MPI application computes a heat distribution matrix. This application models iterative parallel applications that require a large amount

of communication between the boundaries of the MPI processes. The original MPI code was instrumented by inserting the PCM API calls to allow application reconfiguration and checkpointing by the SALSA-MPI framework. The goal of the first experiment was to determine the overhead incurred by the PCM API. As shown in Figure 3, the original and instrumented heat application were run on cluster A with different numbers of nodes. The overhead introduced by the PCM library is reasonable when the number of nodes is small but it increases as the number of nodes increase. This is due to the centralization of the PCM Daemon in the current prototype implementation. In the




**Fig. 3.** The overhead that results from instrumenting the heat distribution program with the PCM API calls on cluster A.

second experiment, the heat distribution program was first run on cluster A. We then increased substantially the load of all the nodes in this cluster by running several computationally intensive programs. We repeated the same experiment, but we migrated the running mpi application to cluster B once the load in cluster A increased. As shown in Figure 4, we notice that the performance of the running mpi application improved once it migrated to a less loaded cluster. The performance improves only when the number of nodes is less than 10 nodes. Again this is due to the limitations of the current prototype implementation. These preliminary results show that reconfigurability and load balancing of parallel mpi applications will improve the performance of the application. The overhead incurred by the PCM library is offset by the application’s overall improved performance. The results also show that having a decentralized architecture is necessary to achieve scalability.

## 5 Discussion and Future Work

We have investigated programming methodologies that promote a separation of the concerns in the implementation of large scientific computations on a large network of computers. High-level programming abstractions provide a natural interface to scientists so that they can concentrate on their domain of expertise. Programming tools map these high-level abstractions into executable units that support efficient communication, dynamic partitioning and load balancing. Run-time middleware infrastructure supports



**Fig. 4.** MPI runs of the Heat Distribution program with and without PCM.

adaptability of executing systems to an evolving underlying network. The presented programming paradigm, languages, and tools are a first step in the unification of parallel and distributed computing by enabling systems to adapt to different and evolving execution environments.

Our initial target applications for SALSA-MPI are parallel adaptive scientific computations. We cannot expect scientists to rewrite or even make significant modifications to extensive libraries of C and C++ software that currently use MPI. The SALSA-MPI architecture allows application programs to run using native C/C++ code, continuing to use MPI for interprocess communication. Applications that provide the minimal required checkpointing functionality can immediately take advantage of SALSA-MPI functionality for dynamic resource allocation, process migration, and fault tolerance. Large-scale computations often already provide checkpointing functionality, so in many cases, this will be a minimal burden. Scientists may choose to provide more fine-grained checkpointing to allow their applications to improve the efficiency of the dynamic migration capabilities, particularly if the computer systems being used are dynamic.

The current prototype implementation is still a work in progress, we are working towards making the SALSA-MPI framework a fully distributed system where the MPI processes or applications reconfiguration are triggered by the IO middleware. The IO middleware should be able to dynamically trigger fine-grain or process migration when the computation to communication ratio is high and coarse-grain or whole application migration when this ratio is low. This will allow the SALSA-MPI framework to accommodate a whole range of scientific and engineering parallel applications.

## References

1. Clark, K., Flaherty, J.E., Shephard, M.S. Appl. Numer. Math., special ed. on Adaptive Methods for Partial Differential Equations **14** (1994)
2. Foster, I., Kesselman, C., Tuecke, S.: The anatomy of the Grid: Enabling scalable virtual organizations. Lecture Notes in Computer Science **2150** (2001) 1–??
3. Foster, I., Kesselman, C., Nick, J., Tuecke, S.: The physiology of the grid: An open grid services architecture for distributed systems integration (2002)

4. Remacle, J.F., Karamete, B., Shephard, M.: Algorithm oriented mesh database. To appear, Proc. 9th Meshing Roundtable, New Orleans (2000)
5. Remacle, J.F., Klaas, O., Flaherty, J.E., Shephard, M.S.: Parallel algorithm oriented mesh database. *Eng. Comput.* **18** (2002) 274–284
6. Bottasso, C.L., Flaherty, J.E., Özturan, C., Shephard, M.S., Szymanski, B.K., Teresco, J.D., Ziantz, L.H.: The quality of partitions produced by an iterative load balancer. In Szymanski, B.K., Sinharoy, B., eds.: Proc. Third Workshop on Languages, Compilers, and Runtime Systems, Troy (1996) 265–277
7. Flaherty, J.E., Loy, R.M., Shephard, M.S., Szymanski, B.K., Teresco, J.D., Ziantz, L.H.: Adaptive local refinement with octree load-balancing for the parallel solution of three-dimensional conservation laws. *J. Parallel Distrib. Comput.* **47** (1997) 139–152
8. Flaherty, J.E., Loy, R.M., Shephard, M.S., Teresco, J.D.: Software for the parallel adaptive solution of conservation laws by discontinuous Galerkin methods. In Cockburn, B., Karniadakis, G., Shu, S.W., eds.: *Discontinuous Galerkin Methods Theory, Computation and Applications*. Volume 11 of *Lecture Notes in Computational Science and Engineering*., Berlin, Springer (2000) 113–124
9. Remacle, J.F., Flaherty, J., Shephard, M.: An adaptive discontinuous Galerkin technique with an orthogonal basis applied to compressible flow problems. *SIAM Review* **45** (2003) 53–72
10. Remacle, J.F., Shephard, M.S.: An algorithm oriented mesh database. *Int. J. Numer. Meth. Engng.* **58** (2003) 349–374
11. Musser, D.R., Saini, A., Stepanov, A.: *STL Tutorial and Reference Guide: C++ Programming With the Standard Template Library*. Addison-Wesley (1996)
12. Gropp, W., Lusk, E., Skjellum, A.: *Using MPI*. M. I. T. Press (1994)
13. Teresco, J.D., Beall, M.W., Flaherty, J.E., Shephard, M.S.: A hierarchical partition model for adaptive finite element computation. *Comput. Methods Appl. Mech. Engrg.* **184** (2000) 269–285
14. Gosling, J., Joy, B., Steele, G.: *The Java Language Specification*. Addison Wesley (1996)
15. Bull, J.M., Smith, L.A., Pottage, L., Freeman, R.: Benchmarking java against c and fortran for scientific applications. In: *Proceedings of ACM Java Grande/ISCOPE Conference*. (2001) 97–105
16. Varela, C., Agha, G.: What after java? from objects to actors. In: *Proceedings of the Seventh International WWW Conference*. Volume 30., *Computer Networks and ISDN Systems: The International Journal of Computer Telecommunications and Networking* (1998) 573–577
17. Varela, C., Agha, G.: Programming dynamically reconfigurable open systems with SALSA. *ACM SIGPLAN Notices. OOPSLA'2001 Intriguing Technology Track Proceedings* **36** (2001) 20–34 <http://www.cs.rpi.edu/~cvarela/oopsla2001.pdf>.
18. Desell, T., ElMaghraoui, K., Varela, C.: Load balancing of autonomous actors over dynamic networks. In: To appear in *Proceedings of the Hawaii International Conference On System Sciences (HICSS-37)*. (2004)
19. Vadhiyar, S.S., Dongarra, J.J.: Srs - a framework for developing malleable and migratable parallel applications for distributed systems (2002)