

Panel on Distributed Simulation and the Grid

Simon J.E. Taylor (Chair)
Centre for Applied Simulation Modeling
Brunel University
Uxbridge, Middlesex, UB8 3PH, UK
E-mail: Simon.Taylor@brunel.ac.uk

George V. Popescu
IBM Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598, USA
E-mail: popescu@us.ibm.com

J. Mark Pullen
C3I Center
George Mason University
Fairfax, VA 22030, USA
E-mail: mpullen@gmu.edu

Stephen J. Turner
Parallel & Distributed Computing Centre
Nanyang Technological University
Nanyang Avenue, Singapore, 639798
E-mail: ASSJTurner@ntu.edu.sg

Abstract

The Grid provides a new and unrivaled technology for large scale distributed simulation as it enables collaboration and the use of distributed computing resources. This panel paper presents the views of four researchers in the area of Distributed Simulation and the Grid. Together we try to identify the main research issues involved in applying Grid technology to distributed simulation and the key future challenges that need to be solved to achieve this goal. Such challenges include not only technical challenges, but also political ones such as management methodology for the Grid and the development of standards. The benefits of the Grid to end-user simulation modelers also are discussed.

1. Introduction

Several members of the community represented by attendees at DS-RT 2004 have been invited to give their position statement on the subject of "Distributed Simulation and the Grid." This is a fascinating topic as it presents new challenges to Grid computing and new opportunities for distributed simulation.

2. Position Statement of Mark Pullen

Wikipedia defines Grid Computing as: "a parallel, distributed system composed of heterogeneous resources located in different places and belonging to

different administrative domains connected over a network using open standards and protocols."

These different places might be opposite sides of the same room or opposite sides of the Earth (say, Singapore and Virginia). Cast this way, I see these challenges:

- "Parallel/distributed computing" is a well-established field, but has not enjoyed the level of focus that it might have. There must be a lot remaining to be learned about effective use of this technology, as it does not appear to have progressed very far beyond the state it held in 1986 when I was managing parallel computing at DARPA. My evaluation is that the funding agencies will need to start taking this problem seriously before progress is made. When even one-tenth as much funding has been applied to parallel/distributed as has already gone into single-processor computing, we might see some maturity and also penetration of the technology into normal application space.
- "Distributed" implies significant network latency; this breaks any fine-grained approach to parallelism. Even some approaches that appear to be coarse-grained (big chunks in different locations) turn out to have a short closed loop among the processes, such that performance is dominated by latency. This problem will be resolved only when the average system architect knows how to use streaming, such that many messages can be processed without feedback from other remote processes.

- “Composed” implies both a discipline of composability that seems to be emerging, and a registry technology (Grid Services) that can most optimistically be described as in its infancy. This area is ripe for major, sustained government funding. It will not be achieved quickly, and industry surely will not make the needed investment until some real achievements beyond the “toy” stage are demonstrated to be working, in a way that makes them economically attractive.
- “Heterogeneous” brings its own set of problems. The Grid people are doing pretty well here, but this one will not be under control until computer manufacturers start using operating systems that implement a Grid suite of open standards. Adding them after-market is inefficient and delays the ability to use the Grid on new computers.
- “Different administrative domains” raises a different kind of problem: management politics. This one will not be overcome until Grid Computing becomes sufficiently established as a discipline that young managers understand well both its potential and its pitfalls, and possess a sufficient “bag of tricks” (management methodology) for dealing with it on a routine basis. This is not very satisfying, I know, but the truth is that Grid Computing will not be mature until it has time enough to mature, with sustained application in the meantime, at least in a widespread academic environment.
- “Connected over a network” is possibly the only part of the Grid Computing definition that is *not* a problem, at least if Grid Computing only requires high performance unicast networking. The Internet continues to evolve improved service rapidly, to the extent that commercial Internet performance is rapidly approaching that of the government's big-bucks Next Generation Internet. While there is a very real chance that the hackers will yet manage to gum up the works, I believe the Internet will continue to grow in size and performance despite their nastiest behaviour.
- Regarding “open standards and protocols,” the sad fact is that developing these is time-consuming and requires long-term support. It does not get done unless somebody with a fair-sized bankroll (typically a large corporation or sometimes a government agency) supports it. Further, to achieve truly effective standards there must be several independent sponsors. The IETF [6] handles this as well as any group I know; their excellent rule is that two independently-developed, interoperating implementations must exist before a standard is

declared. Yet, potentially valuable standards remain unavailable because there is not enough industry interest, or sometimes not enough IESG attention (the IESG is the IETF’s management group) to deal with the issues. I do not envision any improvement here until/unless a significant number of industry players and government agencies recognize they need the Grid standards process to move forward as rapidly as possible (but no faster than that; otherwise you get a half-baked standard).

3. Position Statement of George V. Popescu

The Grid infrastructure defined in [3] provides several components for reliable data management (Grid-FTP), resource allocation and management (GRAM), information discovery and look-up (GIS), authentication and security (GSS), and infrastructure monitoring. The recent OGSA architecture is aimed at virtualizing the Grid infrastructure by defining Grid services abstractions and interactions [4]. By virtualizing the services offered by Grids, new applications such as real-time distributed simulations can be easily deployed on the Grid infrastructure. While Grids provide an extensive set of resource management components, they lack support for some of the network and data management functionality required by distributed simulations.

Data dissemination in interactive applications has multiple goals: dynamic grouping of participants according to their communication interest, efficient data multicasting with guaranteed end-to-end network latency, and reduction of signaling overhead generated by changes in participants’ group membership. Among real-time distributed applications, distributed simulations have specific characteristics: orders of magnitudes smaller network load than streaming applications, stringent end-to-end delay constraints, sensitivity to network delay jitter. In addition, the high dynamics of distributed simulation clients’ group membership require simple multicast group management architectures. This has implication in the design of the application level multicast protocol: the objective of optimal application level multicast (ALM) design is to reduce infrastructure needed to support group communication – data forwarding and control proxies – while satisfying end-to-end network delay constraints.

The proto-typical network architecture of large-scale distributed simulations consists of distributed servers that control data forwarding to dynamic groups of end-system [13]. These are in fact network overlays with star multicast topology where group management

is distributed among a fixed set of control nodes. Such architectures are not efficient from the point of view of network capacity utilization – data distribution load is concentrated at few network controllers – and have limited scalability and do not adapt efficiently to variation in data distribution load generated by group membership dynamics.

The tradeoff between the efficiency of multicast schemes and the amount of signaling for group membership management has large impact in the overall group communication performance. Scaling to a large number of dynamic groups requires multicast protocols with low group membership control load. Among available ALM, few can support efficiently the real-time constraints and group membership dynamics of distributed real-time simulation. In addition, the on-line aspects of efficient data dissemination – group membership and network dynamics – require additional investigation.

When deployed on Internet scale infrastructures, ALM for distributed simulations has multiple objectives: efficient utilization of network resources, satisfying real-time end-to-end delay constraints, minimization of group management control overhead and efficient support for reliable data distribution. In addition, adapting to network and group membership dynamics require self-managing network components, supported through distributed network indexing and index-based routing. Currently Grid infrastructure lack support of application level multicast since the majority of Grid applications need reliable data transfer only, although ongoing research provide the starting point for extending the current Grid infrastructure with real-time communication components [12]. Similar arguments (interoperability, easy deployment of new components) apply for virtualizing network overlay communication to support efficient data dissemination. We define *network virtualization* as an abstraction of network overlay resource management and data distribution services. A required component of *network virtualization* is *distributed network indexing*: a method of abstracting network overlay control and routing by defining a mapping between network overlay nodes and a set of identifiers exposed to the data management component.

The distributed simulation data management component consists of partitioning, load balancing and efficient data indexing and retrieval. Currently available Grid tools (replica management and reliable data transfer – Grid-FTP) provide support for data partitioning and replication only. High data dynamics of distributed simulation require an efficient data indexing and retrieval component. The second

component required to support efficiently distributed simulations on the Grid is the *data management virtualization* component, which abstracts distributed data management (indexing and retrieval) and data transformation operations. Recent examples of network and data indexing techniques are the distributed hash tables (CAN [15], Chord [19]), which manage application data by hashing data identifiers to distributed network overlay nodes. Adaptive distributed indexing networks will support efficiently in the future scalable distributed simulation data management and multi-attribute communication semantics.

Deployment of Internet scale real-time distributed simulations is mainly hampered by the lack of an adaptive communication and data management infrastructure. By providing efficient utilization of network and computing resources, autonomic features [7] could further extend the range of applications deployed on Grids. In particular, the high communication cost of group membership management due to dynamics of client communication interest requires a self-managing distributed infrastructure which can load balance group membership control and data distribution load and adapt to network infrastructure dynamics; adaptive Grids could support more efficiently data dissemination in distributed simulations by reducing the network overlay and group communication control resource requirements.

In conclusion, the currently available Grid components provide adequate support for data replication, and distributed simulation resource allocation and management. Enhancing the grid middleware with a virtualized network and data management component will provide adequate support for distributed simulations on the Grid.

4. Position Statement of Stephen J. Turner

Nowadays, the development of complex modeling and simulation applications usually requires collaborative effort from analysts with different domain knowledge and expertise, possibly at different locations. Furthermore, these simulation systems often require huge computing resources and the data sets required by the simulation may also be geographically distributed. In order to support collaborative model development and to cater for the increasing complexity of such systems, it is necessary to harness distributed resources over the Internet.

In recent years, there has been an explosion of interest in large scale distributed simulation. Much of this activity has centered around the High Level Architecture (HLA) for simulation [2], an IEEE

standard to facilitate interoperability among simulations and promote reuse of simulation models. Using HLA, a large scale distributed simulation can be constructed by linking together a number of geographically distributed simulation models (or federates) into an overall simulation (or federation). However, the HLA does not provide any support for collaborative development of simulation applications, neither does it provide any mechanism for managing the resources where the simulation is being executed.

Grid technology [3] enables collaboration and the use of distributed computing resources, while also facilitating access to geographically distributed data sets. Our vision is a "Grid plug-and-play distributed simulation system", a collaborative distributed simulation environment where analysts at different locations develop, modify, assemble and execute distributed simulations over the Grid [23]. However, a number of important new research challenges need to be addressed before this vision is realized. First, a basic infrastructure providing services to support model discovery and composition is essential in the development of collaborative distributed simulations. Secondly, to conduct simulation experiments easily over distributed resources from different organizations, mechanisms that can provide coordinated and secured simulation executions are required. In addition, to meet the performance requirements demanded by large scale distributed simulations, resource management mechanisms that balance the load and provide fault-tolerance capabilities are needed.

A Grid computing environment consists of a collection of heterogeneous, dynamic, shared resources. These resources may be located at different geographical places and may belong to different administrative domains. The emergence of Grid services [4] and the potential for seamless aggregation, integration and interaction makes it possible to combine computations, experiments, observations, and data to form a powerful simulation environment. Zong et al. [24] describe a preliminary framework for executing HLA-based distributed simulations using Grid services. The RTI control process is managed by an RTI Service and can be dynamically discovered. Simulation models are encapsulated within Federate Services and are assembled through their Grid interface to form a large scale distributed simulation. As different models can be dynamically located, it provides great flexibility. Reusability is inherently provided by the nature of Grid services.

Service composition offers a new and evolving paradigm for building simulation applications. Composition of Grid services allows a number of

components to be combined to achieve the desired computing goals. Service composition involves multi-site interactions, in which the Grid services collaborate with each other in a way defined by the underlying composition mechanism. However, simulation applications are inherently dynamic and heterogeneous and usually require a very large number of service components and very dynamic interactions between these components. Furthermore, the underlying Grid environment where these applications are executing is similarly dynamic and heterogeneous. The combination of the two naturally requires dynamic orchestration and management of service composition.

To provide effective resource management, a number of research issues must be considered. These include resource discovery, federate deployment, load monitoring, dynamic load-balancing, check-pointing and fault-tolerance. We can make use of Grid services to perform the tasks of resource monitoring, coordination of simulation execution and security, while the HLA RTI [2] can be used to perform simulation related tasks such as synchronization and time management. Cai et al. [1] describe a prototype Load Management System (LMS) developed to support the execution of HLA-based simulations over geographically distributed computing resources. Using Grid services, the LMS will match-make the resource requirements of an HLA-based simulation and the resources managed by the resource sharing system, carry out authentication and authorization, schedule the simulation and provide mechanisms for load-balancing and fault-tolerance during the simulation.

Future directions of research include mechanisms to facilitate the discovery, composition and deployment of component simulation models using Grid services. A major challenge is to explore suitable formal approaches to the visual construction, validation and verification of composite simulation applications. Techniques for automating the deployment and execution of such systems (such as mobile agents) are also required. New workflow languages that describe the various component models that constitute the simulation application together with the interactions between them need to be developed.

5. Position Statement of Simon J.E. Taylor

Grid computing offers the potential of conveniently "plugging" in to a secure, scalable parallel and distributed computing resource. There are still many challenges in Grid computing that will in time be addressed by the computer science community. Rather than discussing these in the context of distributed

simulation, I take the position of considering what this emerging technology can offer a group of end user simulation modelers. Before I do this, I will first define what I term the COTS Simulation Package (CSP) simulation modeling community and CSP-based simulation modeling.

It is well known that simulation modeling is a technique that is used to investigate the behaviour of complex, dynamic systems in process-based industries [8] and has associated tools and techniques [17]. In discussions with modelers over the years, a useful common theme in much of this research is the COTS Simulation Package (CSP). Arena, Automod, EPlant, Promodel, Simul8, Taylor and Witness are examples of these (Swain [20] reviews these and many more). These software tools facilitate and limit a range of tasks that a modeler performs in the analysis of discrete event systems found in domains such as manufacturing, health, logistics and commerce. The notion of the CSP tool, and the techniques that it facilitates, has become a useful point of reference [21].

COTS simulation packages are used by simulation modelers mostly for model building, experimentation, animation, visualization and reporting. They have evolved from attempts to build computer environments that support modeling and simulation practice in dynamic process environments. They are typically based on some variant of the discrete event simulation paradigm, i.e. models change state at discrete points in time by scheduled or conditional events and typically represent entities or objects (documents, patients, parts, trains, etc.) in some form that pass through networks of queues and workstations (work queuing at a desk in an office, patients waiting to see a doctor, parts buffered for machining, trains waiting at a station, etc.) Generally, each package has a range of basic model elements (queue, workstation, resource, source, sink, etc.) that are used to build a model via a drag and drop visual interface. Each model element can be modified as is required, either by a menu system or by a package programming language, to better represent the system being studied (for example the queuing logic of a queue or the behaviour of a resource). Entities can be represented and differentiated by attributes. Terminology between packages differs as there is no internationally recognized naming convention.

So what of grid computing and CSP-based simulation modeling? One possibility of course is distributed simulation. The claimed benefits include [14]:

- Team work efficiency is potentially increased, especially in globally distributed teams. This is a typical situation in extended enterprises where

different parts or manufacturing steps are done at different locations.

- The reusability of simulation models and their components is increased. There is no broadly accepted standard for simulation modeling, such as STEP used in product modeling. Within extended enterprises, the use of several simulation tools is common. Distributed simulation allows these models/components to be linked together [18, 11].
- The protection of intellectual property rights (IPR) is facilitated. IPR is an issue if the system being modeled is an external supply chain or a virtual enterprise that includes several legal entities. A single model representing the entire supply chain or virtual enterprise in one company infers that all the modeled knowledge about the other companies is revealed to that company. This procedure is often not acceptable for IPR reasons. Again the “distributed” nature of distributed simulation makes this possible [9, 5].
- In addition to this is the need for group working support, an area that is almost completely ignored and can contribute to major cost savings. The need for regular communication and collaboration throughout the modeling process is generally agreed to be of significant importance [16]. In a study of groupware as exemplified by the net-conference groupware NetMeeting®, it was shown that modelers are very keen to adopt this kind of technology in activities such as remote model debugging and training [22].

Another area that is of interest is experimentation. In a simulation study, a simulation modeler will experiment with his/her model many times to, for example, analyze the sensitivity of model parameters or to optimize some objective function. If the simulation run time is long the depth of experimentation is often reduced due to time constraints on project time. This is further complicated by the fact that models that are stochastic need to be run many times for each experiment to build output statistics. Each run is termed a replication and essentially only differs by the random number seed used in each replicated run.

To summarize, Grid computing can potentially benefit so-called CSP-based simulation modeling by providing a common platform for:

- model interoperability (distributed simulation),
- group working support, and
- high speed distributed experimentation and replication.

However, there is a significant barrier. Nance [10] highlighted the lack of common terminology in simulation modeling. The same is true today. In the light of no standard terminology and a vendor-led community, I ask if it will indeed be possible for the simulation modeling community as described in this position paper to ever benefit from Grid computing. I suggest that it is indeed a worthwhile pursuit as many of the challenges represented by this simulation modeling community could benefit other applications of Grid computing in commerce.

6. Conclusions

This panel paper has presented four views on a new and exciting topic: Distributed Simulation and the Grid. Some of the important research issues involved in applying Grid technology to distributed simulation have been identified. Different views on the key future challenges that need to be solved to achieve this goal have been presented. Although many of these challenges are demanding, both from a technical perspective and also a political one, the potential benefits of Grid computing to end-user simulation modelers are indeed significant.

7. References

- [1] W. Cai, S.J. Turner, and H.F. Zhao, "A Load Management System for Running HLA-based Distributed Simulations over the Grid", *Proc. 6th IEEE International Workshop on Distributed Simulation and Real Time Applications*, 2002, pp. 7-14.
- [2] J. Dahmann, F. Kuhl, and R. Weatherly, "Standards for Simulation: As Simple As Possible But Not Simpler - The High Level Architecture for Simulation", *Simulation*, 71 (6), 1998, pp. 378-387.
- [3] I. Foster, C. Kesselman, and S. Tuecke, "The Anatomy of the Grid: Enabling Scalable Virtual Organizations", *International Journal of Supercomputer Applications*, 15 (3), 2001, pp. 200-222.
- [4] I. Foster, C. Kesselman, J.M. Nick, and S. Tuecke, "The Physiology of the Grid - An Open Grid Services Architecture for Distributed Systems Integration", *Globus Project*, <http://www.globus.org/research/papers/ogsa.pdf>, 2002.
- [5] B.P. Gan, L. Liu, S. Jain, S.J. Turner, W. Cai, and W.-J. Hsu, "Distributed Supply Chain Simulation Across Enterprise Boundaries", *Proc. 2000 Winter Simulation Conference*, 2000, pp. 1245-1251.
- [6] IETF, The Internet Engineering Task Force. <http://www.ietf.org/>
- [7] J. Kephart and D. Chess, "The Vision of Autonomic Computing", *IEEE Computer Magazine*, 36 (1), 2003, pp. 41-50.
- [8] A.M. Law and W.D. Kelton, *Simulation Modeling and Analysis, 3rd edition*, McGraw Hill, New York, 2000.
- [9] K. Mertins and M. Rabe, "Inter-Enterprise Planning of Manufacturing Systems applying Simulation with IPR Protection", *Proc. 5th International Conference on Design of Information Systems for Manufacturing (DIISM)*, Osaka (Japan), 2002, pp. 149-156.
- [10] R.E. Nance, "The Time and State Relationships in Simulation Modeling", *Communications of the ACM*, 24 (4), 1981, pp. 173-179.
- [11] R.J. Paul and S.J.E. Taylor, "What Use is Model Reuse: Is There a Crook at the End of the Rainbow?", *Proc. 2002 Winter Simulation Conference*, 2002, pp. 648-652.
- [12] G. Popescu and Z. Liu, "Stateless Application Level Multicast for Dynamic Group Communication", *Proc. 8th IEEE International Symposium on Distributed Simulation and Real Time Applications*, 2004.
- [13] M. Pullen and D. Wood, "Networking Technology for DIS", *Proceedings of the IEEE*, 83 (8), 1995, pp. 1156-1167.
- [14] M. Rabe, K. Mertins, and S.J.E. Taylor, "Status and Perspectives of Distributed Simulation in Production and Logistics", *Proc. 2004 International IMS Forum*, 2004.
- [15] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker, "A Scalable Content Addressable Network", *Proc. ACM SIGCOMM*, 2001, pp. 161-172.
- [16] S. Robinson and M. Pidd, "Provider and Customer Expectations of Successful Simulation Projects", *Journal of the Operational Research Society*, 3, 1998, pp. 200-209.
- [17] S. Robinson, *Simulation: The Practice of Model Development and Use*, Wiley, Chichester, UK, 2003.
- [18] S. Robinson, R. Nance, R.J. Paul, M. Pidd, and S.J.E. Taylor, "Simulation Model Reuse: Definitions, Benefits and Obstacles", *Simulation Modelling Practice and Theory*, to appear.
- [19] I. Stoica, R. Morris, D. Karger, F. Kaashoek, and H. Balakrishnan, "Chord: A Scalable Content-Addressable Network", *Proc. ACM SIGCOMM 2001 Technical Conference*, San Diego, CA, 2001.
- [20] J.J. Swain, "Power Tools for Visualization and Decision Making: 2001 Simulation Software Survey", *OR/MS Today*, 28 (1), 2001, pp. 52-63.
- [21] S.J.E. Taylor, B.P. Gan, S. Strassburger, A. Verbraeck, "HLA-CSPIF Technical Panel on Distributed Simulation", *Proc. 2003 Winter Simulation Conference*, 2003, pp. 881-887.
- [22] S.J.E. Taylor, S. Robinson, and J. Ladbrook, "An Investigation into the Use of Net-Conferencing Groupware in Simulation Modelling", *Journal of Computers in Information Technology*, to appear.
- [23] G. Theodoropoulos, S.J. Turner, W. Cai and B. Logan, "Large Scale Distributed Simulation on the Grid", *EPSRC e-Science Sister Project GR/S82862/01*, 2003.
- [24] W. Zong, Y. Wang, W. Cai, and S.J. Turner, "Grid Services and Service Discovery for HLA-based Distributed Simulation", *Proc. 8th IEEE International Symposium on Distributed Simulation and Real Time Applications*, 2004.