

International Journal of High Performance Computing Applications

<http://hpc.sagepub.com/>

The Virtual Instrument: Support for Grid-Enabled Mcell Simulations

Henri Casanova, Francine Berman, Thomas Bartol, Erhan Gokcay, Terry Sejnowski, Adam Birnbaum, Jack Dongarra, Michelle Miller, Mark Ellisman, Marcio Faerman, Graziano Obertelli, Rich Wolski, Stuart Pomerantz and Joel Stiles

International Journal of High Performance Computing Applications 2004 18: 3

DOI: 10.1177/1094342004041290

The online version of this article can be found at:

<http://hpc.sagepub.com/content/18/1/3>

Published by:



<http://www.sagepublications.com>

Additional services and information for *International Journal of High Performance Computing Applications* can be found at:

Email Alerts: <http://hpc.sagepub.com/cgi/alerts>

Subscriptions: <http://hpc.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

Citations: <http://hpc.sagepub.com/content/18/1/3.refs.html>

THE VIRTUAL INSTRUMENT: SUPPORT FOR GRID-ENABLED MCELL SIMULATIONS

Henri Casanova¹
Francine Berman¹
Thomas Bartol²
Erhan Gokcay²
Terry Sejnowski²
Adam Birnbaum³
Jack Dongarra⁴
Michelle Miller⁴
Mark Ellisman⁵
Marcio Faerman⁶
Graziano Obertelli⁷
Rich Wolski⁷
Stuart Pomerantz⁸
Joel Stiles⁸

Abstract

Ensembles of widely distributed, heterogeneous resources, or *Grids*, have emerged as popular platforms for large-scale scientific applications. In this paper we present the *Virtual Instrument* project, which provides an integrated application execution environment that enables end-users to run and interact with running scientific simulations on Grids. This work is performed in the specific context of MCell, a computational biology application. While MCell provides the basis for running simulations, its capabilities are currently limited in terms of scale, ease-of-use, and interactivity. These limitations preclude usage scenarios that are critical for scientific advances. Our goal is to create a scientific "Virtual Instrument" from MCell by allowing its users to transparently access Grid resources while being able to steer running simulations. In this paper, we motivate the Virtual Instrument project and discuss a number of relevant issues and accomplishments in the area of Grid software development and application scheduling. We then describe our software design and report on the current implementation. We verify and evaluate our design via experiments with MCell on a real-world Grid testbed.

Key words: grid computing, computational neuroscience

The International Journal of High Performance Computing Applications,
Volume 18, No. 1, Spring 2004, pp. 3–17
DOI: 10.1177/10943420040041290
© 2004 Sage Publications

1 Introduction

Grids (Foster and Kesselman, 1998; Foster et al., 2001) are large collections of resources (computational devices, networks, on-line instruments, storage archives, etc.) distributed over the wide-area; these have enormous aggregate potential and have become popular platforms for running large-scale, resource-intensive applications. Many challenges are to be addressed in order to provide the necessary mechanisms for discovering, accessing, monitoring, and aggregating Grid resources. Consequently, a large effort has been made and is still underway to provide middleware technology as a base Grid software infrastructure (Foster and Kesselman, 1999; Foster et al., 2001, 2002). However, although middleware provides fundamental building blocks it is not designed to be used directly by Grid users. Instead, Grid application-level tools must be provided with the goal of both building new and higher-level functionality on top of base Grid services, as well as hiding the complexity of the Grid from the end-user. One approach is to provide programming models that implement high-level abstractions for building Grid applications (Benyon et al., 2001; van Nieuwpoort et al., 2001; Nakada et al., 2002; Karonis et al., 2003), or even general purpose Grid application development environments (Berman et al., 2001). Another approach is to implement execution environments in which a user can "drop" his/her application for Grid execution while maintaining the convenience and the illusion of a desktop execution. Such environments include Grid portals, which have been implemented successfully for many scientific applications and provide users with a familiar Web browser interface to launch and monitors application runs on Grid resources: for instant the GAMESS portal, <https://gridport.npaci.edu/GAMESS/>, and the Telescience portal, <https://gridport.npaci.edu/Tel->

¹ SAN DIEGO SUPERCOMPUTER CENTER AND DEPT. OF COMPUTER SCIENCE AND ENGINEERING, UNIVERSITY OF CALIFORNIA, SAN DIEGO

² COMPUTATIONAL NEUROBIOLOGY LABORATORY, SALK INSTITUTE

³ SAN DIEGO SUPERCOMPUTER CENTER

⁴ DEPT. OF COMPUTER SCIENCE, UNIVERSITY OF TENNESSEE, KNOXVILLE

⁵ NATIONAL CENTER FOR MICROSCOPY AND IMAGING RESEARCH, UNIVERSITY OF CALIFORNIA, SAN DIEGO

⁶ DEPT. OF COMPUTER SCIENCE AND ENGINEERING, UNIVERSITY OF CALIFORNIA, SAN DIEGO

⁷ DEPT. OF COMPUTER SCIENCE, UNIVERSITY OF CALIFORNIA, SANTA BARBARA

⁸ PITTSBURGH SUPERCOMPUTER CENTER

escience/ (Allen et al., 2001). As a result, several efforts have provided toolkits to help in the development of Grid portals; see Thomas et al. (2001), Novotny (2002), Suzumura et al. (2002) and the Grid Portal Collaboration (<http://www.ipg.nasa.gov>). Another solution is to build integrated software environments targeted to specific applications or classes of applications. The work described in this paper belongs in this last category; our ultimate goal is to build software that makes an application behave as a scientific *Virtual Instrument* (VI) that the user can easily configure, observe, and dynamically control. We note that our work could be integrated as part of a Grid portal with minimal effort.

This work is performed in the specific context of the MCell application; see the MCell webpages at the Pittsburgh Supercomputer Center, <http://www.mcell.psc.edu>, and at the Salk Institute, <http://www.mcell.cnl.salk.edu> (Stiles et al., 1996, 1998). MCell is a computational biology simulation framework that is used by neuroscientists to study diffusion and chemical reactions of molecules in living organisms. A single execution instance of MCell consists of multiple simulations, each producing output that is then analyzed by neuroscientists en masse. Although MCell provides the core functionality for running simulations, its capabilities are currently limited in terms of scale, ease of use, and interactivity. For example, scientists often observe interesting phenomena that emerge in the middle of an MCell run. If MCell could be redirected to concentrate on these phenomena while executing, a great deal of time could be saved. This and other limitations preclude usage scenarios that are critical for scientific advances. The goal of the VI project is to alleviate most of the limitations of MCell usage and to provide an integrated Grid application execution environment for MCell users. This environment should provide transparency for access to the Grid as well as computational steering capabilities.

In this paper we describe our accomplishments in terms of software design and development for the Grid. This development primarily entails the realization of the VI software, a complete run-time system that supports steerable MCell executions. We present experimental results obtained when running the application on a Grid testbed and draw conclusions about the effectiveness of the VI implementation and design.

This paper is organized as follows. In Section 2 we introduce MCell and highlight specific limitations of its current usage scenarios. In Section 3 we motivate the VI project, highlight issues of application scheduling and Grid software development, and describe the VI software design and implementation, which we verify experimentally and discuss in Section 4. Section 5 discusses related work and Section 6 concludes the paper with future directions.

2 Molecular Biology Simulations with MCell

2.1 MCELL OVERVIEW

MCell (see the MCell webpages at the Pittsburgh Supercomputer Center, <http://www.mcell.psc.edu>, and at the Salk Institute, <http://www.mcell.cnl.salk.edu>; Stiles et al., 1996, 1998) uses Monte Carlo algorithms to simulate simultaneous diffusion and chemical reactions of molecules in complex three-dimensional (3D) spaces. Highly realistic reconstructions of cellular or subcellular boundaries can be used to define 3D diffusion spaces, which can then be populated with different molecules (Stiles et al., 2001). Such molecules might react with others that are released periodically from different locations within the structure, to simulate the production of biological signals. The diffusing molecules move according to a 3D random walk based on a Brownian motion model. Possible reaction events, such as binding and unbinding, are tested on a molecule-by-molecule basis using random numbers and Monte Carlo probability values. The advantages and significance of this approach are detailed in Stiles and Bartol (2001).

In essence, computational modeling with MCell encompasses four steps, each of which can require considerable computing resources.

- (i) **Surface design or reconstruction.** In simple cases, a set of planes might be used to define diffusion boundaries. In complex cases, cell membranes can be reconstructed as tessellated meshes from electron microscope data, and may contain on the order of 10^6 triangles.
- (ii) **Model visualization and design.** Different types of molecules must be added to the surfaces and spaces according to realistic biological distributions and densities. The total number of molecules is highly variable but can easily reach or exceed 10^6 even for a surface area or reaction volume much smaller than a single cell.
- (iii) **Simulation.** This step involves repeatedly running MCell with varied input parameters and Monte Carlo random number streams. The total number of such runs can range from 10^2 to 10^5 and beyond. We detail relevant usage scenarios for this step in the next section.
- (iv) **Visualization and analysis of results.** In the simplest case this might require two-dimensional (2D) plotting of one output parameter as a function of time. In the more typical case, some combination of 2D plotting and 3D imaging and/or animation is required to visualize the simulation's output.

At present, all simulation objects and run-time conditions are specified using a high-level model description lan-

guage (MDL) designed for readability by scientists. When a simulation is run, one or more MDL input files are parsed to create the simulation objects, and then execution begins for a specified number of time-step iterations. MCell is highly optimized for speed and for memory usage. This makes it possible to run individual simulations of complex structures on single processors rather than using parallel architectures.

So far, MCell simulations have been used to study synapses, structures used by nerve cells to communicate with themselves and other cells. MCell (and its predecessors) originally focused on the popular vertebrate neuromuscular junction, the synapse between a nerve cell and a muscle cell (Salpeter, 1987; Stiles et al., 2001). MCell's Monte Carlo simulations have been successfully employed to obtain a variety of new results (Bartol et al., 1991; Anglister et al., 1994; Stiles et al., 1996, 1998, 1999, 2001). In addition, MCell has been in limited release (see the MCell webpages at the Pittsburgh Supercomputer Center, <http://www.mcell.psc.edu>, and at the Salk Institute, <http://www.mcell.cnl.salk.edu>) to a worldwide group (~25) of Neuroscience and other research laboratories since 1997 (Egelman and Montague, 1998; Egelman et al., 1998; Rao-Mirotnik et al., 1998; Gieger et al., 1999). MCell is currently the object of many development efforts and current simulations are allowing scientists to explore new areas of cellular physiology.

2.2 MCELL USAGE SCENARIOS

Since MCell models are now approaching the level of structural and biochemical complexity present in living cells, the models typically contain numerous input parameters that can be varied independently. Consequently, simulations can span an enormous range of computational and data requirements. We detail here three relevant usage scenarios. We give orders of magnitude for the aggregate simulation CPU time assuming a single 2.0 GHz Pentium processor.

- (A) **“Look & See”.** A small number of MCell runs are used to determine the predicted behavior of the modeled system under limited input conditions; between one hour and several days of CPU time are required.
- (B) **Parameter Fitting.** Tens to thousands of runs may be required to identify input parameter values which produce model output that matches given criteria such as experimental measurements; several weeks of CPU time.
- (C) **Parameter Sweep.** The scale of individual simulations is similar to that for the parameter fitting scenario, but many thousands of runs are required to map a region of the input parameter space; any-

where from one month of CPU time to several years or decades.

Even though (A) has been the most common scenario in the early stages of the MCell project, it is increasingly being replaced/complemented by scenarios (B) and (C). These last two scenarios require a tremendous amount of compute and storage resources. For (B), the user generally navigates toward a “best fit” by iterative parameter adjustments made according to some potentially ad hoc heuristics. Thus, a high degree of interactivity between the user and the computing resources is desirable to maximize productivity. Scenario (C) does not require interactivity as the user has already identified an “interesting” region of the parameter space to explore, perhaps via scenario (B). In Casanova et al. (2001), we gave an example of a small-scale simulation for (C), which required approximately three months of aggregate CPU time and generated 94 GB of raw output, which were then reduced to 600 KB of synthesized output. Note that in all scenarios the CPU time of individual simulations can vary by several orders of magnitude solely depending on input parameter values.

2.3 CURRENT LIMITATIONS

Currently, MCell imposes severe limitations on the usage scenarios described in the previous section. Ideally, users would have access to integrated software which guides them through the four steps identified in Section 2.1 and which enables all three usage scenarios on large-scale distributed computing environments.

In its current incarnation, MCell consists of a single executable which takes MDL files as input. The user is responsible for creating these files and for managing each MCell “project” in an ad hoc fashion. The user is entirely responsible for running the individual simulations and collecting the output. This involves labor-intensive activities such as resource selection, remote process creation/monitoring, fault-detection and restart, or application data movements. These tasks are generally performed via a set of ad hoc scripts. In scenarios (B) and (C), this proves to be infeasible for most users given the desired scale of the simulations. In addition, there is no support for interactive simulation as required in scenario (B).

The MCell executable generally produces one or more output files. Users are responsible for averaging, post-processing, visualizing, and analyzing output files. In scenarios (B) or (C), this amounts to manipulating and mining large datasets, again in an ad hoc fashion. MCell users typically employ the file system as a database for application data, which does not scale and cannot support scenarios (B) or (C) adequately. Generally, OpenDX

(<http://www.opendx.org>) is used for most rendering and visualization tasks.

Our earlier work on the AppLeS Parameter Sweep Template (APST; Casanova and Berman, 2003) provides limited support for (C) in terms of running MCell simulation on Grid resources. It is a clear improvement over the traditional usage described above. However, since APST is general purpose, it fails to address most of the MCell-specific limitations listed above. As scenarios (B) and (C) are the future of MCell simulations for new scientific discoveries, it is critical to provide corresponding comprehensive software support. This is the overall goal of the VI project.

3 The Virtual Instrument

In this section we describe the VI project in terms of specific goals, relevant research issues, and software.

3.1 GOALS

The main motivation behind the VI software development effort is to address the limitations highlighted in Section 2.3. More specifically, this is accomplished by providing the following capabilities.

Framework for MCell project development The VI must provide a framework in which MCell users can easily specify individual “projects” in terms of input parameters, initial ranges for those parameters, number of repeats for Monte Carlo averaging, nature of output files, and nature of output post-processing steps (see Section 3.2 for a detailed description of MCell projects). The only component of an MCell simulation that cannot be automatically created is the core MDL code as it embodies the user’s conceptual model. The VI must provide a framework for users to plug in their core MDL code and be freed of all other responsibilities. This framework can easily be embedded as part of a user interface. Finally, the VI should handle all application data management issues. This can be done, for instance, with a relational database.

User interface At the moment, MCell does not provide any user interface to facilitate MCell project instantiation and management. The VI should provide a graphical interface for users to create MCell projects within the framework described above. In addition, that interface should be able to invoke data visualization and rendering capabilities provided by tools like OpenDX (<http://www.opendx.org>). A full-fledged user interface for MCell is an intensive development project and is not our focus here. Instead, we aim at providing a simple interface that will allow us to explore computer science research issues involved when running large-scale distributed MCell simulations in scenarios (B) and (C).

Transparent deployment The VI should handle the logistics of application deployment on behalf of the user. This includes resource discovery, authentication/authorization, remote job creation/control, application data movements, fault-detection and recovery. This can be achieved by building on the base Grid software infrastructure.

Interactive simulation In order for scenario (B) to be effective, the VI must provide a way for users to interact with running simulations in order to guide, or *steer*, the computation. Users must be able to direct the search away from certain regions of the parameter space to be explored, and to concentrate on other regions, based on real-time intermediate application results. This requires that the VI allow the creation and cancellation of application tasks on the fly.

High performance Given the scale of MCell simulations in scenarios (B) and (C), it is critical that the VI exploit available resource effectively. This is to be achieved by the use of scheduling strategies, and can build on our previous work (Casanova et al., 2000a). However, in this work, there is the added complexity of computational steering; how does one schedule (and reschedule) an application whose computational goals are constantly being changed and/or refined by the user? Our goal is to develop resource allocation strategies that reduce application execution time (e.g. search time) in the presence of user steering.

3.2 MCELL PROJECTS IN THE VI FRAMEWORK

Before describing our work on scheduling and on Grid software design and development, we introduce the notion of an *MCell project*, which can be created and executed by the VI user. The structure of a project is depicted in Figure 1 and consists of: (i) a set of parameter specifications (number of parameters, data types, initial value ranges); (ii) a set of MDL scripts written by the user – the MDL core; (iii) potential additional input files such as large geometry files that have been produced by 3D reconstruction of electron microscope data. The MCell simulation consists of a (generally large) number of parameter space point evaluations, or meta-tasks (n meta-tasks are shown in the figure). Each evaluation consists of an instantiation of the parameter values and of a number of identical MCell tasks, each using different streams of random numbers for Monte Carlo simulation. Each task produces raw output files that are then averaged and synthesized into final output. This final output is typically orders of magnitude smaller (in terms of bytes) than raw and averaged output. These final data must be analyzed by the user to understand the behaviors

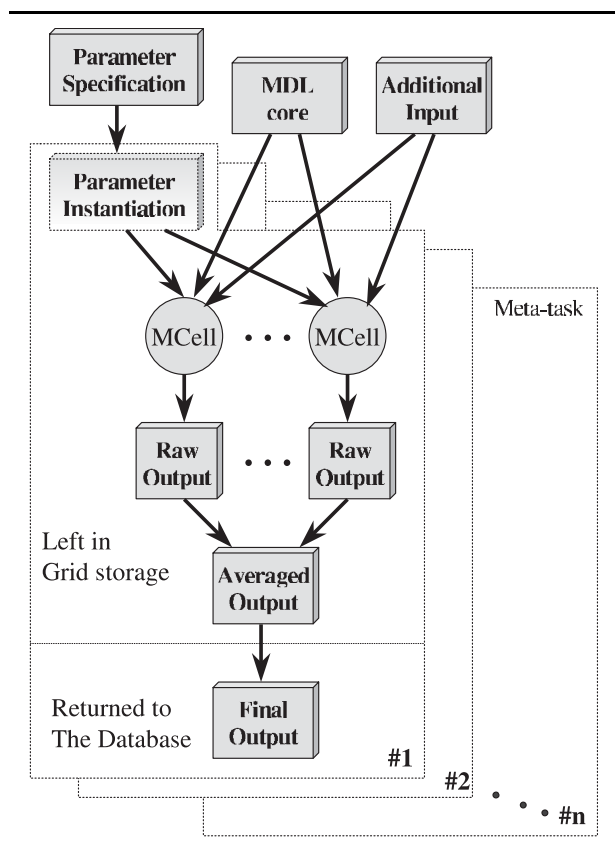


Fig. 1 The structure of an MCell project in the VI Framework. A project consists of n meta-tasks, and each meta-task consists of a number of identical MCell tasks whose outputs are averaged and synthesized into final output data.

of the simulated biological system. This must be done on-the-fly to steer the computation in scenario (B). Note that raw and average output data are still of interest to the MCell user and may need to be retrieved in order to perform in-depth analyses. Note also that the amount of computation required for each meta-task depends on the parameter instantiations. In other words, some meta-tasks require more computation than some others (typically by up to three orders of magnitudes).

A run of an MCell project consists in executing large numbers of independent meta-tasks and generating both intermediate and final output.

3.3 GRID COMPUTING ISSUES

The two main issues that we have explored in the VI project are that of Grid application development and application scheduling. We briefly discuss these issues below.

3.3.1 Resource Access and Data Management on the Grid.

The VI project faces most of the issues inherent in Grid computing as it seeks to make the use of the Grid as transparent as possible so that the user can focus on the MCell simulation rather than on the logistics of application deployment. To that end, we reuse most of the available Grid infrastructure to achieve automatic resource discovery (Czajkowski et al., 2001), resource access (Casanova and Dongarra, 1997; Czajkowski et al., 1998; Basney and Livny, 1999), security (Foster et al., 1998), distributed data management (Chervenak et al., 2000), and resource monitoring (Wolski, 1997). Our work on the VI builds on our experience when developing the APST project (Casanova and Berman, 2003). APST provides a simple, generic way to run parameter sweep applications and is currently used by MCell users for scenario (C).

One of the lessons we learned with APST is that targeting several underlying technologies for deploying user application makes it possible to (i) gain early acceptance from the users and (ii) increase the number of resources available to applications. This is the case because Grid computing is still an emerging technology and is not yet ubiquitous. Although the Open Grid Software Architecture (OGSA) standard (Foster et al., 2002) is rapidly gaining momentum, our goal is to enable MCell users to run simulations on their resources today.

Consequently, the VI targets a number of Grid services, which can be used simultaneously to expand the range of resources available to a single MCell simulation. In addition, the VI provides default mechanisms that use SSH to start remote jobs and move application data. SSH does not provide the levels of job control and the scalability offered by, say, Globus (Foster and Kesselman, 1999). However, our experience with APST is that users generally start using SSH mechanisms and progressively move towards Grid middleware technology as their simulation needs grow in scale. The main notion here is that current Grid application execution environments should be able to use whatever Grid middleware is available, but also degrade to default ubiquitous mechanisms if necessary. We expect this design to evolve to pure OGSA when the standard becomes more ubiquitous.

Given the life-span of MCell simulations, it is critical that the core VI software be resilient to software and hardware crashes. In addition, the VI should automatically handle all application data management on behalf of the user. Consequently, we use a relational database in order to maintain persistent state about running MCell projects, data generated by those projects, and available resources. This database has two roles. First, it allows the VI software to be resilient to faults; all state is periodically saved into the database and can be used for restart. Secondly, it provides a structure for storing, retrieving,

and mining application data, which is fundamental for achieving the first goal in Section 3.1. Our approach is to store only final application output data in the database (see Figure 1). Raw and intermediate output, which can be enormous, is left in place on remote Grid storage resources and can be downloaded on demand by the user.

3.3.2 Application Scheduling. A research direction explored in the VI project is that of application scheduling in a computational steering context. Scheduling sets of non-identical, independent tasks onto sets of distributed, heterogeneous resources has long been identified as an NP-hard problem (Ullman, 1975). Therefore, much research work has been dedicated to the development of appropriate scheduling heuristics; see Braun et al. (1999) for a survey. Grid computing adds several challenges to the traditional scheduling problem; resources are not only heterogeneous, they exhibit dynamic performance behaviors due to sharing among users. Also, they are located on diverse network topologies interconnected over the wide-area. To address these issues, *adaptive scheduling* has been employed with success (Berman et al., 1996; Berman, 1998). As a result, we developed adaptive scheduling strategies for MCell in our earlier work (Casanova et al., 2000a), focusing on scheduling data movement, data staging, and data duplication, with respect to storage and compute resource locations and characteristics.

In this work, we have to cope with the complexity added by computational steering, that is the problem of scheduling an application whose computational goals change over time according to potentially arbitrary user behaviors. Computational steering is a difficult problem that has been addressed by several researchers (Jablonowski et al., 1993; Vetter and Schwan, 1995, 1997; Geist et al., 1997; Parker et al., 1998). These efforts mostly addressed the problems of consistency of state among components of tightly coupled applications. In the limited context of MCell, consistency is not a key issue as the application consists of large sets of tasks which can be stopped and restarted independently, with little need for synchronization. Therefore, we focus on a resource allocation strategy that takes steering into account solely to achieve high performance.

The VI user typically employs the following general search strategy. The user's goal is to locate some particular point in the parameter space that satisfies some subjective criteria. An initial set of parameter space points (uniformly) distributed over the parameter space are selected for computation. As results come back from the VI, they are displayed to the user who can then assign *levels of importance* to regions of the parameter space. Regions with higher levels of importance are more promising and should therefore be completed sooner. This can be achieved by assigning appropriate fractions of the avail-

able compute resources to the exploration of each region. For instance, if the user has identified three regions that should be explored and assigned levels of importance 2, 2, and 1, then each of the first two regions should get 40% of the resources, and the last region should get 20%. The key idea is that rather than ordering regions by their level of importance and exploring the most promising regions first, it is more efficient to explore all regions concurrently but at different rates, to avoid being trapped in local optima in regions that were initially the most promising. Our contribution is that we reason at the resource allocation level rather than at the algorithmic level, which makes our approach applicable to a wide variety of steering behaviors and search algorithms (including both interactive user steering and search algorithms).

Our approach is to assign a priority to each point of the parameter space, corresponding to the level of importance of the region to which the point belongs. If each of the current n points being computed has a priority p_i , then point i should get $p_i / \sum_{j=1}^n p_j$ percent of the available resources. In Faerman et al. (2002) we have proposed and evaluated a number of strategies for computing priorities and for scheduling computational tasks according to these priorities. We refer the reader to that paper for all details. Our main result was that the use of priorities for allowing concurrent region explorations leads to significant performance improvements and we have quantified how priorities should be computed, given levels of importance that users assigned to regions.

3.4 VIRTUAL INSTRUMENT SOFTWARE

The VI software follows a strict object-oriented design and is constructed of three principal components: a software daemon to manage resources and remotely run jobs; a user interface to allow users to initiate, run, monitor, and stop MCell projects; and a database to store final application results and user-entered data. These components can run on separate machines.

Figure 2 depicts the interactions of the three main components of the VI architecture: the Daemon, the Interface, and the Database. The Daemon interacts with resources via Grid services. These services allow the Daemon to discover resources, start and control remote jobs, move data between distributed storage locations, and monitor resources as well as the running application. The Daemon uses the Database to store information such as the available resources, the user-defined specifications of running MCell projects, and the status of these running projects, including their pending tasks. To the greatest extent, the Daemon uses an out-of-core approach, so that if it fails, the relevant information about running MCell projects is in the Database. The only application data stored in the Database are MCell final output that can be visualized and analyzed

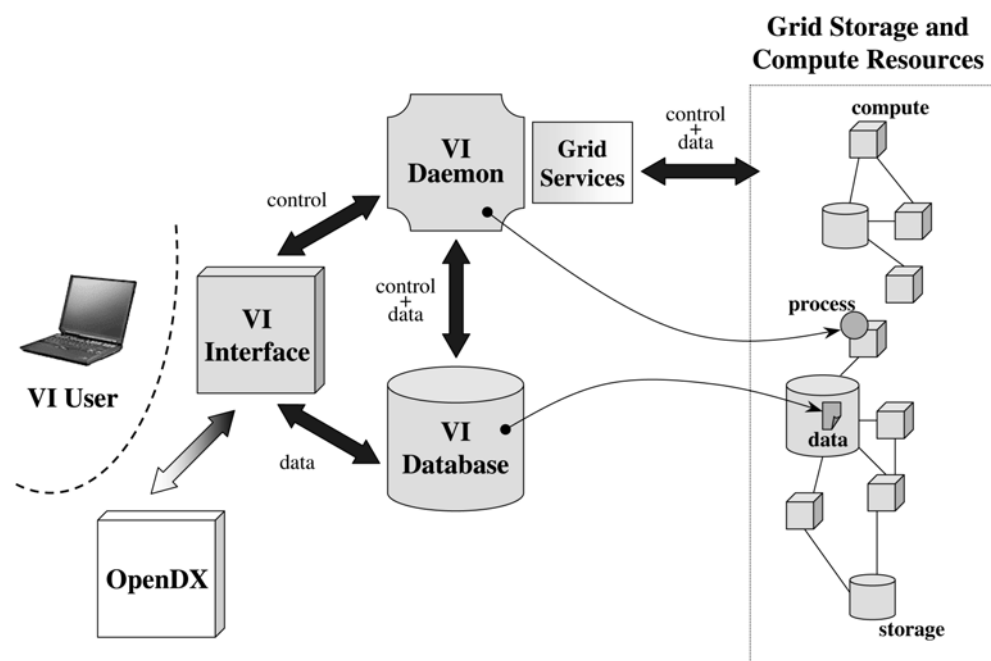


Fig. 2 The VI architecture. The three main components are the Daemon, the Database, and the Interface.

by the user and used to steer further simulations. As depicted in Figure 2, the Database also keeps track of the location of all raw and intermediate data, which are left in place in Grid storage until explicitly retrieved by the user (see Section 3.3).

The Interface allows the user to steer the computation and to perform visualization. The Interface communicates control information to the Daemon, including commands to create, start, steer, and stop MCell projects or trigger the retrieval of a particular output file or final output data. Visualization of the data can be performed by OpenDX at the user's direction, as invoked from the Interface.

The main responsibility of the Daemon is to schedule and actuate file transfers and computations using available computational and network resources. These functions are performed by three classes within the Daemon: the Project class, the Scheduler class, and the Actuator class. The Project keeps track of all of the parameter space points and task inter-dependences. For example, in Figure 1, it is the Project that is aware of the requirement to complete several runs of MCell with their parameter instantiations before running a post-processing task to average the output. The Scheduler retrieves information on tasks from the Project, sets their relative priorities (see

Section 3.3), and assigns tasks to resources accordingly. The Scheduler is designed as a base class so that alternate scheduling strategies can be easily integrated as they are developed. After tasks have been assigned to resources by the Scheduler, the Actuator launches them on Grid resources. As with the Scheduler, the Actuator is designed as a base class, permitting specialization for various remote job execution and data transfer methods from various Grid middleware services.

The use of a relational database has several advantages. It makes the design of the Daemon more simple in terms of data structures, and makes it possible to recover from failures. In addition, the Interface does not need to implement an ad hoc protocol with the Daemon, but can just pull data out of the Database in a standard fashion. In particular, it is possible to make complex SQL queries to mine application data. Furthermore, the use of a separate database allows users to start an MCell project, disconnect, and check the status of the simulation from any location.

3.5 STATUS OF THE IMPLEMENTATION

At the moment, the VI software consists of approximately 20,000 lines of C++, using the Standard Template

Library and pthreads. We opted for MySQL to implement the Database as it is well accepted by the Linux community and provides a standard API. In the current release, the actuators within the Daemon target SSH, Globus's GRAM (Czajkowski et al., 1998), NetSolve (Casanova and Dongarra, 1997), and batch schedulers (the portable batch system, <http://www.openpbs.com>; IBM Corporation, 1993) for starting/monitoring remote jobs, scp and GridFTP for moving application data on the Grid. The VI interface is still underway and at the moment we provide several interfaces. First, we have implemented a text-only interface for evaluation purposes. This interface allows us to gather information about user behaviors and requirements for converging towards a graphical interface. This interface is also written in C++ on top of VI components. We have also implemented a Web-based portal to the VI so that users can check on progress and perform simple data-mining tasks on the final application output. In addition, we have implemented a stand-alone Java-based GUI which is being used as the base for the final VI GUI. At the moment it implements the same functionality as the portal. Finally, we have implemented a simulator in order to evaluate our scheduling/steering strategies (see Section 3.3). The simulator is written with SIMGRID (Casanova, 2001; Legrand et al., 2003), a toolkit specially designed for the study of application scheduling in distributed computing systems, and has been integrated with the VI software. This allows us to simulate a variety of user behaviors and to test and validate the VI implementation throughout development.

An alpha version of the VI software was released to a limited number of MCell users/developers in February 2002 for evaluation. The software was subsequently enhanced and hardened and demonstrated at the SC'02 conference. The demonstration involved a simulation consisting of 11,360 MCell tasks and was executed on a testbed aggregating diverse computing resources in Japan and California (including batch-scheduled MPPs, batch-scheduled clusters, interactive clusters, and individual workstations). The simulation was launched and monitored by a user on the conference floor in Baltimore, Maryland. A beta version of the software was released to MCell users in February 2003.

The software, information about installation, and further details about the implementation can be found on the project's Webpage at http://grail.sdsc.edu/projects/vi_itr.

4 Experiments

We have performed a large run of an MCell project on a Grid platform over the course of several days. Note that these experiments did not include any user steering as our goal here is solely to verify and evaluate the VI soft-

ware design. The software supports steering and we refer the reader to Faerman et al. (2002) for a discussion of steering and scheduling.

4.1 THE MCELL SIMULATION

In this section we give background for and describe the MCell simulation that we ran in our experiments. This simulation is called *r_disk* and targets a synapse. A synapse is a highly organized cellular structure that forms at the narrow junction between two neurons (or between a neuron and a muscle cell) in the nervous system. Chemical communication occurs across this synaptic cleft through a process called synaptic transmission. In synaptic transmission neurotransmitter molecules are released through the membrane of an excited pre-synaptic neuron and activate specific neurotransmitter receptor molecules on the membrane surface of a post-synaptic cell. Receptor activation results in a transient signaling event which might lead to excitation of the post-synaptic cell.

The *r_disk* MCell simulation attempts to map the parameter space of synaptic transmission for a canonical, simple model of synaptic geometry and assumes acetylcholine (ACh) and acetylcholine receptor (AChR) as the neurotransmitter/receptor system. The synaptic geometry includes infinite planes for the pre- and post-synaptic membranes separated by a synaptic cleft distance fixed at 20 nm. There is a disk-shaped patch of receptors on the post-synaptic membrane. A fixed number of 10,000 ACh molecules are released from a point centered over the receptor patch. In this model we explore a seven-dimensional parameter space which includes the radius of the receptor patch, the diffusion constant of ACh, and five additional chemical kinetics parameters describing the various on rates, off rates, and conformation-change rates for the reaction mechanism between ACh and AChR. The overall computation time for this simulation amounts to approximately 329 days of computation on a 1.5 GHz Pentium III.

4.2 THE TESTBED

The testbed for our experiments consists of resources at three sites: the Tokyo Institute of Technology (TITECH) in Japan, the San Diego Supercomputer Center (SDSC) and the Department of Computer Science and Engineering (CSE), both at the University of California, San Diego (UCSD). All resources run some version of the Linux operating system. Table 1 summarizes overall characteristics of the resources and how they are accessed. Two of the sites, TITECH and CSE, provide interactive resources. Resources at TITECH can be reserved via an e-mail reservation system, and thus were not time-shared with other

Table 1
Testbed for the experiments presented in Section 4.

| Site | CPUs | Access |
|--------|---|-------------------------|
| TITECH | Presto-III cluster 20 dual-cpu Athlon, 1.59 GHz | Ssh |
| SDSC | Meteor cluster 93 dual-cpu PIII, 730 MHz to 996 MHz | PBS |
| CSE | GRAIL workstations 12 Athlon and PIII, 200 MHz to 1.5 GHz | 4 via Ssh, 8 via Globus |

applications for our run. The resources at CSE were not reserved and were thus time-shared with local user jobs. The resources at SDSC are accessed via the Portable Batch System (<http://www.openpbs.com>) batch queuing system. In these runs, the VI software used four simultaneous submissions to the PBS queue, with requests for 4, 8, 16, and 32 processors. The Daemon and Database, as well as the user, are located on one of the CSE machines. At the end of the execution all synthesized application output is available in the Database and all intermediate application output is accessible in remote storage.

4.3 RESULTS

Figure 3 depicts the cumulative amount of computation performed with respect to time. We show a curve for the total platform (Total) as well as a breakdown per site (SDSC, TITECH, and CSE). Note that we do not plot the cumulative number of tasks completed. Indeed, this number is misleading as MCell tasks vary widely in computational costs. For this simulation, the amount of computation required varied by as much as a factor of 200 between MCell tasks. Instead, we ran an off-line benchmark of the MCell tasks involved in this simulation on reference CPUs (identical nodes of the Presto-III cluster). This allowed us to associate each task in the simulation with a “relative computational cost” (i.e. the amount of time the task requires on the reference CPU). This provides the basis for the cumulative computational cost plotted on the y-axis of Figure 3. An alternate approach would have been to use a performance model, but no accurate performance model for MCell tasks was available at the time of these experiments. In addition, the computational cost estimates obtained via the benchmark were used to estimate wallclock time requests passed to the PBS batch system on the Meteor cluster.

We make three main observations on the data presented in Figure 3.

- (i) All sites contribute to the computation. As expected, the site contributing the most to the computation is SDSC. Overall, SDSC contributed to 83% of the overall computation, TITECH to 13%, and CSE to 4%. The entire simulation (which amounts to approximately 329 days of computa-




Fig. 3 Cumulative amount of computation performed during the experimental run on all sites and per site.

tion on a 1.5 GHz Pentium III) was completed in a little over 150 h.

- (ii) All curves level off before the end of the execution, due to the typical “waiting for the last task” syndrome. When scheduling a set of independent tasks, it is generally a good idea to schedule long tasks first (Hummel et al., 1996; Hagerup, 1997). However, as discussed earlier, we do not have a performance model for MCell tasks. In other words, there is no way for the VI scheduler to determine automatically which tasks will be long (or short). As a result, it is likely that a number of long tasks will be scheduled at the end of the execution, leading to the leveling of the curves in Figure 3. A possible approach would be to run the aforementioned off-line benchmark and have the scheduler use the benchmark timings to make scheduling decisions. However, this is time- and resource-consuming, and it would need to be done for each new MCell project. More importantly, note that in a real usage scenario the user would continually steer the simulation and add new tasks. Therefore, reducing the execution time of a

fixed number of tasks is not the focus of this work (unlike that of our previous work (Casanova et al., 2000a, 2000b)). In this work we focus on overall throughput and thus do not consider the “waiting for the last task” problem to be significant.

- (iii) The curve for SDSC exhibits more of a step pattern than the curve for TITECH or CSE. This is due to the use of the PBS batch system at SDSC; MCell tasks are submitted in batches, each batch spends time waiting in the queue, and thus tasks in batches tend to return results in bursts. By contrast, the TITECH resources are available in interactive mode and thus complete tasks at a relatively steady rate. This raises an interesting question regarding the use of batch-scheduled resources for running applications that consist of many independent, small tasks, as well as for running interactive applications. How many simultaneous requests and which request sizes should be sent to a batch-scheduled resource? In the current VI software these values can be fixed arbitrarily (as seen in Section 4.2). We leave the investigation of how these values should be chosen to maximize throughput as future work.

Figure 4 plots the cumulative number of failures that occurred during the entire execution for the SDSC and CSE resources. We do not show a curve for the TITECH site as no failure was experienced during this experiment on TITECH resources. 60 failures occurred during the run, meaning that on average the VI experienced approximately 0.6 failures per hour. These failures were caused by several factors, including actual resource downtimes, network time-outs, and software failures (e.g. SSH disconnections). The VI software detects failures, attempts a number of retries, and then marks resources as “failed” for a fixed amount of time before attempting other retries. Therefore, some of the failures plotted in Figure 4 correspond to several occurrences of a failure for a single resource.

4.4 DISCUSSION

The experimental results obtained with the first VI implementation provide convincing evidence the following.

- (i) The VI software is functional and makes it possible to support large MCell runs on large-scale Grid testbeds. Furthermore, the software shields the user from the logistics of application deployment and provides the illusion of a desktop execution.
- (ii) The VI software design is effective. The use of an out-of-core relational database for storing all




Fig. 4 Cumulative number of failures during the experimental run for SDSC and CSE.

information pertaining to the resources and the application proved to be scalable. In these experiments and others that are not reported in this paper, the use of the database was never a bottleneck and scales with the size of the application and the platform. In fact, the main bottleneck for application deployment is the overhead for launching remote computations on distant Grid resources.

- (iii) The VI implementation is resilient to resource faults and failures. Our experiments show that failures occur in Grid environments at a non-negligible rate and due to several causes. The VI software appropriately detects failures and attempts a number of retries. This is critical for long-running applications such as MCell.

While a number of features and capabilities can be added to the VI software, in particular to the user interface, it provides sufficient functionality and robustness to enable the new generation of MCell simulations. Furthermore, all logistics of application deployment are hidden to the user and application data are automatically managed on behalf of the user. We conclude that the current version of the software meets the goals outlined in Section 3.1.

5 Related Work

Our work is related to a number of large efforts that seek to provide Grid application execution environments for scientific simulations: GriPhyN (<http://www.griphyn.org>); National Virtual Collaboratory for Earthquake Engineering Research (<http://www.neesgrid.org>); [47] Particle

Physics Data Grid (<http://www.ppdg.net>). Related works also include portal activities (Thomas et al., 2001; Novotny, 2002) and the VI software could ultimately be integrated as a user portal.

Computational steering has been an active field of research and several projects have provided models, methodologies, and software for steering scientific applications: SCIRUN (Parker et al., 1998), VASE (Jablonowski et al., 1993), Progress (Vetter and Schwan, 1995), Magellan (Vetter and Schwan, 1997), CUMULVS (Geist et al., 1997). One of the main challenges addressed in these works is the notion of state consistency. Several techniques from the area of distributed systems and fault-tolerance have been used successfully to build high performance consistent computational steering environments. Our work is related to those efforts in that we provide computational steering capabilities. However, given the structure of MCell simulations, i.e. parallel searches with loose task and data synchronization requirements, state consistency is not a crucial issue. Therefore, our work focuses mostly on performance issues and proposes a scheduling/steering strategy based on task priorities for appropriate resource sharing.

This work builds on our earlier work on the APST (Casanova and Berman, 2003), which is related to projects such as Nimrod (Abramson et al., 2000) or ILAB (Yarrow et al., 2000). APST provides a generic Grid application execution environment for Parameter Sweep Applications. These applications consist of large numbers of computational tasks that exhibit few or no inter-dependences. This category of applications encompasses many methodologies such as Monte Carlo simulations, parametric studies, and parameter searches, and arises in many fields of science and engineering. This work uses APST as a learning experience to provide a full-fledged execution environment customized for MCell. APST addresses a few of the limitations listed in Section 2.3 and is currently used for medium-scale MCell parameter sweep runs. Neither APST, Nimrod, nor ILAB provide capabilities for computational steering.

6 Future Work and Conclusions

In this paper we have presented the VI project, which targets the deployment of large-scale, interactive MCell simulations. MCell is a molecular biology simulator, which has gained great popularity in the computational neuroscience community. Although the current MCell software provides basic capabilities to run simulations, it does not enable interactive simulation, and leaves many responsibilities to the user in terms of deployment, scheduling, and data management. These limitations preclude the use of MCell for large-scale executions, espe-

cially on the Grid platform. The goal of the VI project is to provide an integrated Grid execution environment for MCell that offers interactive computational steering capabilities. We have described contributions of our Grid software development effort and have given a brief account of our work in the area of application scheduling. We have then described the VI software in detail and presented validation experiments conducted for a real-world MCell application on a multi-site Grid testbed. These results are conclusive as they show that the VI software is functional and the design scalable.

Many future directions are currently being explored in this project. A beta version of the VI software was recently released to MCell users. Our ultimate goal is to deploy the software in a production environment: (i) to further validate our implementation; (ii) to log information about usage and learn about user behaviors; (iii) to enable new disciplinary results. Ultimately, the Virtual Instrument will have a large and quantifiable impact on the MCell community. It will lead to new scientific advances that would not be possible without the Grid platform and without our fully integrated software environment.

ACKNOWLEDGMENTS

The authors wish to acknowledge the San Diego Supercomputer Center and the Matsuoka Laboratory at the Tokyo Institute of Technology for providing computational resources that made this work possible. This work was supported by the National Science Foundation under Award ACI-0086092.

BIOGRAPHIES

Dr. Henri Casanova is an adjunct Professor of Computer Science and Engineering at the University of California, San Diego, a Research Scientist at the San Diego Supercomputer Center, and the founder and director of the Grid Research And Development Laboratory (GRAIL) at UCSD. His research interests are in the area of parallel, distributed, Grid, and Internet computing with emphases on modeling, scheduling, and simulation. Dr. Casanova obtained his B.S. from the Ecole Nationale Supérieure d'Electronique, d'Electrotechnique, d'Informatique et d'Hydraulique de Toulouse, France in 1993, his M.S. from the Université Paul Sabatier, Toulouse, France in 1994, and his Ph.D. from the University of Tennessee, Knoxville in 1998.

Thomas M. Bartol, Jr. is a senior research associate in the Computational Neurobiology Laboratory at the Salk Institute for Biological Studies. He received his Ph.D. in cellular and molecular neurobiology from Cornell Uni-

versity in the laboratory of Dr. Miriam Salpeter using electrophysiology and computational techniques to study neuromuscular synaptic transmission. His thesis studies laid the stepping stones that led to the creation of MCell, a general Monte Carlo simulator of cellular microphysiology, which he coauthored with Joel Stiles. He has published research articles and book chapters and has presented several tutorial workshops on MCell and its use in the study of synaptic transmission and cellular physiology. His current research centers on the further development of large-scale, realistic computer simulation methods to study pre- and post-synaptic components of synaptic transmission and synaptic plasticity.

Dr. Francine Berman is Director of the San Diego Supercomputer Center and the National Partnership for Advanced Computational Infrastructure, ACM Fellow, Professor of Computer Science and Engineering at the University of California, San Diego, and founder of the UCSD Grid Laboratory. Over the last two decades, her research has focused on the development of software, tools and models for Parallel and more recently Grid Computing environments. Dr. Berman has served on numerous program committees, advisory boards and steering committees and currently serves as one of the two Principal Investigators for NSF's TeraGrid project and as co-PI on the Extensible Terascale Facility (ETF) project. Dr. Berman obtained her B.A. from the University of California, Los Angeles in 1973, her M.S. from the University of Washington, 1976, and her Ph.D. from the University of Washington in 1979.

Adam Birnbaum is the Grid Applications Lead at San Diego Supercomputer Center, University of California, San Diego. He served as a research programmer at GRAIL, UCSD during the research reported here. In previous work, Mr. Birnbaum has worked as a software developer and manager in a variety of startup companies, including Critical Concepts, Eight Cylinder Studios, and On2.com. Mr. Birnbaum obtained an A.B. from Harvard College in 1990, and an M.A. in Physics from SUNY Stony Brook in 1992.

Jack Dongarra received a Bachelor of Science in Mathematics from Chicago State University in 1972 and a Master of Science in Computer Science from the Illinois Institute of Technology in 1973. He received his Ph.D. in Applied Mathematics from the University of New Mexico in 1980. He worked at the Argonne National Laboratory until 1989, becoming a senior scientist. He now holds an appointment as University Distinguished Professor of Computer Science in the Computer Science Department at the University of Tennessee and is an Distinguished R&D Participant in the Computer

Science and Mathematics Division at Oak Ridge National Laboratory (ORNL), and an Adjunct Professor in the Computer Science Department at Rice University. He specializes in numerical algorithms in linear algebra, parallel computing, use of advanced-computer architectures, programming methodology, and tools for parallel computers. He is a Fellow of the AAAS, ACM, and the IEEE and a member of the National Academy of Engineering.

Mark H. Ellisman is a professor of Neuroscience and Bioengineering and the Director of the Center for Research in Biological Structure at UCSD. Prof. Ellisman directs the National Center for Microscopy and Imaging Research (NCMIR), an internationally acclaimed technology development center and research resource established by the National Institute of Health (NIH). He has received numerous awards including a Jacob Javits award from the NIH and the Creativity Award from the National Science Foundation, and he is a Founding Fellow of the American Institute of Biomedical Engineering. His scientific contributions include work on basic molecular and cellular mechanisms of the nervous system and development of advanced technologies in microscopy and computational biology. He is a pioneer in the development of three-dimensional light and electron microscopy and combined application of these image acquisition tools and computational technologies to achieve greater understanding of the structure and function of the nervous system. His group was the first to introduce the idea of "Telemicroscopy" by demonstrating the network-enabled remote use and sharing of a high-energy electron microscope in 1992 and then developed practical systems now in use by researchers in the US and abroad.

Dr. Marcio Faerman is a researcher at the San Diego Supercomputer Center and was recently a member of GRAIL at UCSD. His research interests include distributed data intensive computing, grid computing, application-level scheduling, computer networks, and performance modeling and prediction. Dr. Faerman obtained his B.S. in 1993, his M.S. in 1995 in Electrical Engineering from Universidade Federal do Rio de Janeiro and his Ph.D. in Computer Science from UCSD in 2003.

Erhan Gokcay received his B.S. and M.S. in Electrical and Electronic Engineering from the Middle East Technical University in Ankara, Turkey, in 1986 and 1991 respectively, and his Ph.D. in Computer Science and Engineering from the University of Florida, Gainesville in 2000. He worked as an application programmer for ASELNAN Military Electronics and Telecommunica-

tions Company, in Ankara, Turkey, and as a system programmer in STFA Enercom Computer Center, in Ankara, Turkey, from 1986 to 1990. He worked as a technical manager at Tulip Computers in Ankara, Turkey, until 1991. He worked as a network administrator and hardware supervisor in Bilkent University, in Ankara, Turkey, from 1991 to 1993. Currently, he is working as a system manager at the UCSD Swartz Center and as a scientific programmer at the Salk Institute, in San Diego, California.

Michelle Miller received a B.S. in Anthropology in 1984 and an M.E. in Computer Science in 1998 from the University of Utah. She worked in the Innovative Computing Laboratory at the University of Tennessee from 1999 until 2003 as a research programmer. She currently works at the National Optical Astronomy Observatory designing and developing software to process and manage astronomical data. Her research interests include high performance distributed computing, scientific computing, operating systems, component-based software system construction and software fault tolerance.

Graziano Obertelli received his Laurea degree from the Dipartimento di Scienze dell'informazione at the Università di Milano. He is currently a research programmer in the Department of Computer Science at University of California, Santa Barbara. His research interests are in the area of Grid computing.

Stuart M. Pomerantz received a B.A. in English and B.A. in Psychology from Penn State University in 1990 and a B.S. in Mathematics from the University of Pittsburgh in 1995. From 1995 to 2000 he worked as a Systems Analyst at the University of Pittsburgh Department of Mathematics operating the departmental computing facilities. In 2000 he received an M.S. in Information Science from the University of Pittsburgh. He joined the Pittsburgh Supercomputing Center in December 2000 as a Research Programmer applying his expertise in graphics programming and software design to the Visible Human and MCell projects.

Terrence Sejnowski is an Investigator with the Howard Hughes Medical Institute, a Professor at the Salk Institute for Biological Studies, where he directs the Computational Neurobiology Laboratory, and a Professor of Biology at UCSD, where he is Director of the Institute for Neural Computation. He founded Neural Computation, the leading journal in neural networks and computational neuroscience, and is also the President of the Neural Information Processing Systems Foundation, which oversees the annual NIPS Conference.

Joel R. Stiles, M.D., Ph.D. (University of Kansas) is Associate Professor in the Mellon College of Science and Pittsburgh Supercomputing Center at Carnegie Mellon University. He also holds adjunct faculty positions at the University of Pittsburgh in the Department of Neuroscience and Center for Computational Biology and Bioinformatics. His research interests focus primarily on quantitative modeling of synaptic transmission, with a particular emphasis on neuromuscular transmission and pathology. More generally, his interests include Monte Carlo and hybrid simulation approaches to 3D reaction-diffusion problems, multiscale modeling, surface and volume meshing for subcellular and cellular reconstructions, and large-scale 2D and stereoscopic rendering technology. He is one of the original co-authors of MCell, a Monte Carlo simulator of cellular microphysiology, and is also the principal author of DReAMM (Design, Render, and Animate MCell Models).

Rich Wolski is an assistant professor of Computer Science at the University of California, Santa Barbara. His research interests include Grid computing, distributed computing, scheduling, and resource allocation. In addition to AppLeS, he leads the Network Weather Service project which focuses on on-line prediction of resource performance. He has developed EveryWare, a set of tools for portable Grid programming. He is also leading the G-commerce project studying computational economies for the Grid. He received his B.S. in Computer Science from the California Polytechnic State University at San Luis Obispo, and his M.S. and Ph.D. degrees from the University of California Davis, LLNL campus.

REFERENCES

- Abramson, J., Giddy, D., and Kotler, L. May 2000. High performance parametric modeling with Nimrod/G: Killer application for the global Grid? In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS)*, Cancun, Mexico, pp. 520–528.
- Allen, G., Shalf, J., Benger, W., Dramlitsch, T., Goodale, T., Hege, H.-C., Lanfermann, G., Merzky, A., Radke, T., and Seidel, E. 2001. Cactus tools for Grid applications. *Cluster Computing* 4(3):179–188.
- Anglister, L., Stiles, J. R., and Salpeter, M. M. 1994. Acetylcholinesterase density and turnover number at frog neuromuscular junctions, with modeling of their role in synaptic function. *Neuron* 12:783–794.
- Bartol, T. M., Land, B. R., Salpeter, E. E., and Salpeter, M. M. 1991. Monte Carlo simulation of miniature endplate current generation in the vertebrate neuromuscular junction. *Biophysical Journal* 59(6):1290–1307.
- Basney, J. and Livny, M. 1999. Deploying a high throughput computing cluster. In *High Performance Cluster Computing*, Vol. 1, Chap. 5, Prentice Hall, Englewood Cliffs, NJ.

- Berman, F. 1998. *The Grid, Blueprint for a New computing Infrastructure*, I. Foster and C. Kesselman, editors, Chap. 12. Morgan Kaufmann Publishers, San Mateo, CA.
- Berman, F., Wolski, R., Figueira, S., Schopf, J., and Shao, G. 1996. Application-level scheduling on distributed heterogeneous networks. In *Proceedings of SC'96*, Pittsburgh, PA.
- Berman, F., Chien, A., Cooper, K., Dongarra, J., Foster, I., Dennis Gannon, L. J., Kennedy, K., Kesselman, C., Reed, D., Torczon, L., and Wolski, R. 2001. The GrADS project: Software support for high-level grid application development. *International Journal of High Performance Computing Applications* 15(4):327–344.
- Beynon, M., Kurc, T., Catalyurek, U., Chang, C., Sussman, A., and Saltz, J. 2001. Distributed processing of very large datasets with datacutter. *Parallel Computing* 27(11):1457–1478.
- Braun, R., Siegel, H., Beck, N., Boloni, L., Maheswaran, M., Reuther, A., Robertson, J., Theys, M., Yao, B., Hensgen, D., and Freund, R. April 1999. A comparison study of static mapping heuristics for a class of meta-tasks on heterogeneous computing systems. In *Proceedings of the 8th Heterogeneous Computing Workshop (HCW'99)*, San Juan, Porto Rico, pp. 15–29.
- Casanova, H. May 2001. Simgrid: a toolkit for the simulation of application scheduling. In *Proceedings of the IEEE International Symposium on Cluster Computing and the Grid (CCGrid'01)*, Brisbane, Australia, pp. 430–437.
- Casanova, H. and Berman, F. 2003. Parameter sweeps on the Grid with APST. In *Grid Computing: Making the Global Infrastructure a Reality*, Chap. 33, Wiley, New York.
- Casanova, H. and Dongarra, J. 1997. NetSolve: a network server for solving computational science problems. *International Journal of Supercomputer Applications and High Performance Computing* 11(3):212–223.
- Casanova, H., Legrand, A., Zagorodnov, D., and Berman, F. May 2000a. Heuristics for scheduling parameter sweep applications in grid environments. In *Proceedings of the 9th Heterogeneous Computing Workshop (HCW'00)*, Cancun, Mexico, pp. 349–363.
- Casanova, H., Obertelli, G., Berman, H., and Wolski, R. November 2000b. The AppLeS parameter sweep template: user-level middleware for the Grid. In *Proceedings of SC'00*, Dallas, TX.
- Casanova, H., Bartol, T., Stiles, J., and Berman, F. 2001. Distributing MCell simulations on the Grid. *International Journal of High Performance Computing Applications* 14(3):243–257.
- Chervenak, A., Foster, I., Kesselman, C., Salisbury, C., and Tuecke, S. 2000. The data Grid: towards an architecture for the distributed management and analysis of large scientific datasets. *Journal of Network and Computer Applications* 23(3):187–200.
- Czajkowski, K., Foster, I., Karonis, N., Kesselman, C., Martin, S., Smith, W., and Tuecke, S. 1998. A resource management architecture for metacomputing systems. In *Proceedings of IPPS/SPDP'98 Workshop on Job Scheduling Strategies for Parallel Processing*, Orlando, FL.
- Czajkowski, C., Fitzgerald, S., Foster, I., and Kesselman, C. August 2001. Grid information services for distributed resource sharing. In *Proceedings of the 10th IEEE Symposium on High-Performance Distributed Computing*, San Francisco, CA.
- Egelman, D. and Montague, P. 1998. Computational properties of peri-dendritic calcium fluctuations. *Journal of Neuroscience* 18(21):8580–8589.
- Egelman, D., King, R., and Montague, P. 1998. Interaction of nitric oxide and external calcium fluctuations: a possible mechanism for rapid information retrieval. *Progress in Brain Research* 118:199–211.
- Faerman, M., Birnbaum, A., Casanova, H., and Berman, F. November 2002. Resource allocation for steerable parallel parameter searches. In *Proceedings of the Grid Computing Workshop*, Baltimore, MD, pp. 157–169.
- Foster, I. and Kesselman, C., editors. 1998. *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann Publishers, San Francisco, CA.
- Foster, I. and Kesselman, C. 1999. Globus: A toolkit-based Grid architecture. In *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann Publishers, San Francisco, CA, pp. 259–278.
- Foster, I., Kesselman, C., Tsudik, G., and Tuecke, S. 1998. A security architecture for computational grids. In *Proceedings of the 5th ACM Conference on Computer and Communications Security*, San Francisco, CA, pp. 83–92.
- Foster, I., Kesselman, C., and Tuecke, S. 2001. The anatomy of the Grid: enabling scalable virtual organizations. *International Journal of High Performance Computing Applications* 15(3):200–222.
- Foster, I., Kesselman, J., Nick, J., and Tuecke, S. June 2002. The physiology of the Grid: an open grid services architecture for distributed systems integration. *Open Grid Service Infrastructure WG, Global Grid Forum*.
- Geist, G., Kohl, J., and Papadopoulos, P. 1997. CUMULVS: providing fault tolerance, visualization, and steering of parallel applications. *International Journal of Supercomputer Applications and High Performance Computing* 11(3):224–235.
- Gieger, J., Roth, A., Taskin, B., and Jonas, P. 1999. Glutamate-mediated synaptic excitation of cortical interneurons. In *Handbook of Experimental Pharmacology, Retinoids, Ionotropic glutamate receptors in the CNS*, Vol. 141, P. Jonas and H. Moyner, editors, Springer-Verlag, Berlin, pp. 363–398.
- Hagerup, T. 1997. Allocating independent tasks to parallel processors: an experimental study. *Journal of Parallel and Distributed Computing* 47:185–197.
- Hummel, S. F., Schmidt, J., Uma, R. N., and Wein, J. June 1996. Load-sharing in heterogeneous systems via weighted factoring. In *Proceedings of the 8th Annual ACM Symposium on Parallel Algorithms and Architectures*, Padua, Italy, pp. 318–328.
- IBM Corporation. 1993. *IBM LoadLeveler User's Guide*.
- Jablonski, D., Bruner, J., Bliss, B., and Haber, R. 1993. VASE: the visualization and application steering environment. In *Proceedings of Supercomputing 1993*, Portland, OR, pp. 560–569.
- Karonis, N., Toonen, B., and Foster, I. 2003. MPICH-G2: a grid-enabled implementation of the message passing interface. *Journal of Parallel and Distributed Computing* 63(5):551–563.

- Legrand, A., Marchal, L., and Casanova, H. May 2003. Scheduling distributed applications: the SimGrid simulation framework. In *Proceedings of the 3rd IEEE International Symposium on Cluster Computing and the Grid (CCGrid'03)*, Tokyo, Japan.
- Nakada, H., Matsuoka, S., Seymour, K., Dongarra, J., Lee, C., and Casanova, H. November 2002. An overview of GridRPC: a remote procedure call API for Grid computing. In *Proceedings of the Grid Computing Workshop*, Baltimore, MD, pp. 274–279.
- Novotny, J. 2002. The Grid portal development kit. *Concurrency and Computation: Practice and Experience* 14(13–15):1129–1144.
- Parker, S., Miller, M., Hansen, C., and Johnson, C. January 1998. An integrated problem solving environment: the SCIRun computational steering system. In *Proceedings of the 31st Hawaii International Conference on System Sciences (HICSS-31)*, Kohala Coast, Hawaii, Vol. VII, pp. 147–156.
- Rao-Mirotnik, R., Buchsbaum, G., and Sterling, P. 1998. Transmitter concentration at a three-dimensional synapse. *Journal of Neurophysiology* 80(6):3163–3172.
- Salpeter, M. M., editor. 1987. *The Vertebrate Neuromuscular Junction*, Alan R. Liss, New York, pages 1–54.
- Stiles, J. R. and Bartol, T. M. 2001. Monte Carlo methods for simulating realistic synaptic microphysiology using MCell. In *Computational Neuroscience: Realistic Modeling for Experimentalists*, E. DeSchutter, editor, CRC Press, Boca Raton, FL.
- Stiles, J. R., Van Helden, D., Bartol, T. M., Salpeter, E. E., and Salpeter, M. M. 1996. Miniature endplate current rise times < 100 μ s from improved dual recordings can be modeled with passive acetylcholine diffusion from a synaptic vesicle. In *Proceedings of the National Academy of Sciences USA* 93:5747–5752.
- Stiles, J. R., Bartol, T. M., Salpeter, E. E., and Salpeter, M. M. 1998. Monte Carlo simulation of neurotransmitter release using MCell, a general simulator of cellular physiological processes. In *Computational Neuroscience*, J. M. Bower, editor, Plenum Press, New York, pp. 279–284.
- Stiles, J. R., Kovyazina, I. V., Salpeter, E. E., and Salpeter, M. M. 1999. The temperature sensitivity of miniature endplate currents is mostly governed by channel gating: evidence from optimized recordings and Monte Carlo simulations. *Biophysical Journal* 77:1177–1187.
- Stiles, J. R., Bartol, T. M., Salpeter, M. M., Salpeter, E. E., and Sejnowski, T. J. 2001. Synaptic variability: new insights from reconstructions and Monte Carlo simulations with MCell. In *Synapses*, W. M. Cowan, T. C. Südhof, and C. F. Stevens, editors, Johns Hopkins University Press, Baltimore, MD, pp. 681–731.
- Suzumura, T., Nakada, H., Saito, M., Matsuoka, S., Tanaka, Y., and Sekiguchi, S. November 2002. The Ninf portal: an automatic generation tool for the grid portals. In *Proceedings of Java Grande*, Seattle, WA, pp. 1–7.
- Thomas, M., Mock, S., Boisseau, J., Dahan, M., Mueller, K., and Sutton, S. August 2001. The GridPort toolkit architecture for building grid portals. In *Proceedings of the 10th IEEE International Symposium on High Performance Distributed Computing (HPDC-10)*, San Francisco, CA.
- Ullman, J. 1975. NP-complete scheduling problems. *Journal of Computer and System Sciences* 10:434–439.
- van Nieuwpoort, R., Kielmann, T., and Bal, H. June 2001. Efficient load balancing for wide-area divide-and-conquer applications. In *Proceedings of the 8th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'01)*, Snowbird, UT, pp. 34–43.
- Vetter, J. and Schwan, K. 1995. PROGRESS: a toolkit for interactive program steering. In *Proceedings of the 1995 International Conference on Parallel Processing*, Urbana-Champaign, IL, pp. 139–149.
- Vetter, J. and Schwan, K. 1997. High performance computational steering of physical simulations. In *Proceedings of IPPS'97*, Geneva, Switzerland, pp. 128–132.
- Wolski, R. August 1997. Dynamically forecasting network performance using the network weather service. In *6th High Performance Distributed Computing Conference*, Portland, OR, pp. 316–325.
- Yarrow, M., McCann, K., Biswas, R., and Van der Wijngaart, R. December 2000. An advanced user interface approach for complex parameter study process specification on the information power Grid. In *GRID 2000*, Bangalore, India.