TeraGrid Enabled National Scale Epidemiological Models for the Models of Infectious Disease Agent Study (MIDAS)

Diglio A. Simoni, Douglas Roberts, Eric Solano
RTI International
Research Triangle Park, North Carolina, 27709, USA
{dsimoni, droberts, solano}@rti.org

Madhav Marathe, Stephen Eubank, Keith Bissett, Chris Barrett, Karla Atkins, Chen Jiangzhuo
Virginia Bioinformatics Institute
Blacksburg, Virginia, USA
{mmarathe, seubank, kbisset, cbarrett, katkins, chenj}@vbi.vt.edu

Abstract

Distributed memory agent-based models (ABMs) have been used in recent years to predict the spread of disease in human populations. A design for a Grid version of an epidemiological ABM is presented. This project is currently underway within the Models of Infectious Disease Agent Study (MIDAS) research network, and is a collaboration between Virginia Bioinformatics Institute (VBI) and RTI International (RTI). Important design elements of this project include the partitioning of input data sets, global I/O, and synchronization of distributed computational resources. The computational resources used on this project are funded by a TeraGrid Medium Resource Allocation Committee (MRAC) grant of 200,000 system units (SUs), funded over a two-year period. The project was granted Advanced Support for TeraGrid Applications (ASTA) status because of its importance to modeling national epidemiological scenarios.

Keywords: Infectious Diseases, Mathematical Models, Epidemiology, High Performance Computing, High Level Programming, Clusters, Networking Distributed Memory Models, Agent Based Models, TeraGrid, Message Passing Interface.

Resumen

Los modelos matemáticos de memoria distribuida y basados en individuos (ABM, por sus siglas en inglés) han sido utilizados en años recientes para predecir la dispersión de enfermedades en poblaciones humanas. En este trabajo se presenta un diseño en versión “Grid” o red de un modelo epidemiológico ABM. Este proyecto se está desarrollando actualmente como parte del esfuerzo del grupo de análisis de modelos matemáticos para el estudio de mecanismos de dispersión de enfermedades infecciosas (MIDAS, por sus siglas en inglés). Este proyecto es un esfuerzo de colaboración entre el Instituto de Bioinformática de Virginia (VBI) y RTI International (RTI). Algunos elementos de diseño importantes de este proyecto incluyen la partición de bases de datos de entrada, I/O global, y la sincronización de recursos computacionales distribuidos. Los recursos computacionales usados en este proyecto son financiados a través de una concesión del Comité de Distribución de Recursos de la Red “TeraGrid” (MRAC, por sus siglas en inglés) y consiste en 200,000 unidades de sistema (SUs), los cuales son financiados por un periodo de dos años. Al proyecto se le otorgó estatus de “Apoyo Avanzado para Aplicaciones de TeraGrid (ASTA)” debido a su importancia en el desarrollo de escenarios epidemiológicos nacionales en los EEUU.

Palabras Clave: Enfermedades Infecciosas, Modelos Matemáticos, Epidemiología, Computación de Alto Rendimiento, Programación de Alto Nivel, Conglomerados de Computadores, Redes, Modelos de Memoria Distribuida, Modelos Basados en Individuos, TeraGrid, Interfase de Transmisión de Mensajes.

1 Introduction

The Models of Infectious Disease Agent Study (MIDAS) is a research partnership between the United States National Institutes of Health (NIH) and the scientific community whose purpose is to develop computational models for policymakers, public health workers, and other researchers to assist them in making better-informed decisions about emerging infectious diseases, both man-made and naturally-occurring. MIDAS researchers are working to develop models that may assist the public health community understand how best to respond during outbreaks and epidemics. The Virginia Bioinformatics Institute (VBI) is one of the original three MIDAS research groups, and is the current home of EpiSims, a distributed discrete-event agent-based epidemiological simulation system. An important part of the process of using the TeraGrid is to map data to the computers on the TeraGrid. Figure 1 shows a representation of a gridded population that must be created before being applied to a problem on the TeraGrid. It shows the computational resources of the TeraGrid, a translation of the gridded population.

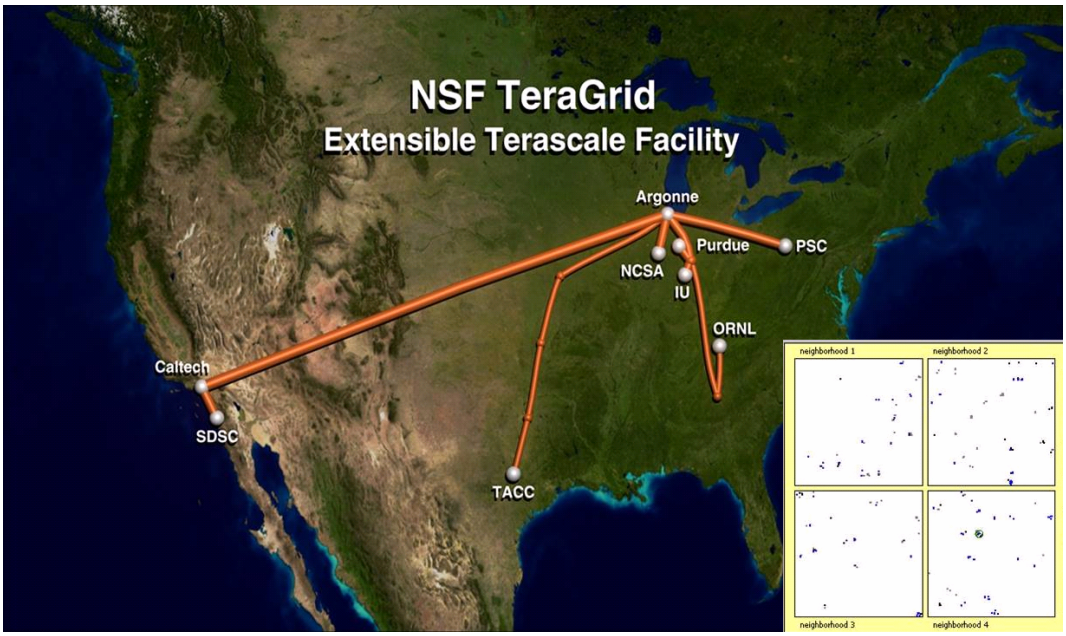

Figure 1. Mapping Social Networks to the TeraGrid

2 MIDAS Research Groups and Scientific Activity

The Research Groups focus on models that address research questions that span many aspects of spatiotemporal/biological issues. These include host-pathogen relationships, disease epidemiology, disease surveillance methods and pandemic response strategies. The Research Groups focus on information-driven research rather than hypothesis-driven investigations. MIDAS model developers use real or simulated data that are widely available through the MIDAS web site (www.epimodels.org). The Informatics Group developed and maintains the web site, and manages the information and data resources that modelers may access. The MIDAS Research groups actively publish their results in prestigious peer reviewed journals (see [3] through [6]).

The Research Groups have done extensive work in both infectious diseases transmitted primarily by the respiratory route (influenza) and vector-borne diseases (malaria, dengue). Future MIDAS Research may involve the combination of Agent/Individual Based Models or Equation Based Models with ecological niche modeling to predict spatial dynamics of vector insects and human cases of vector borne diseases such as dengue and malaria. Some non-MIDAS research has used ecological niche modeling to study vector-borne diseases ([7] and [8]) while other efforts concentrate on the application of individual based models to analyze vector borne diseases ([1] and [2]).

3 TeraGrid Enabled Models

In this research project, the "end users" are the members of the MIDAS Research Groups. The product will consist of a set of technical reports on the performance characteristics of agent based simulations when they are run on grid architectures.

The challenges building TeraGrid Models include considering heterogeneity and temporal variations, determining the scale and scope of simulations, computing speed and memory requirements and maintaining efficiency and sustained performance. Some important design elements for a grid enabled model include partitioning and distributing the problem, activity resolution, synchronization of local versus global requirements, global I/O, use of Grid-Enabled Implementation of the Message Passing Interface (MPICH-G2) and use of General Parallel File System (GPFS). Figure 2 shows the hardware used in the VBI/RTI Medium Resource Allocation (MRAC).

4 Synchronization Issues

The project involves a distributed, discrete-event ABM architecture. There are two distinct synchronization requirements: on-cluster node-node synchronization (very tight) and inter-cluster synchronization (loose). On-cluster node synchronization requirements are determined by the intra-region population mobility patterns. Initially, we will be putting one city on a TeraGrid cluster and other cities on other TeraGrid cluster. Inter-cluster synchronization requirements are determined by inter-region travel requirements, where each cluster is responsible for modeling a population region (city).

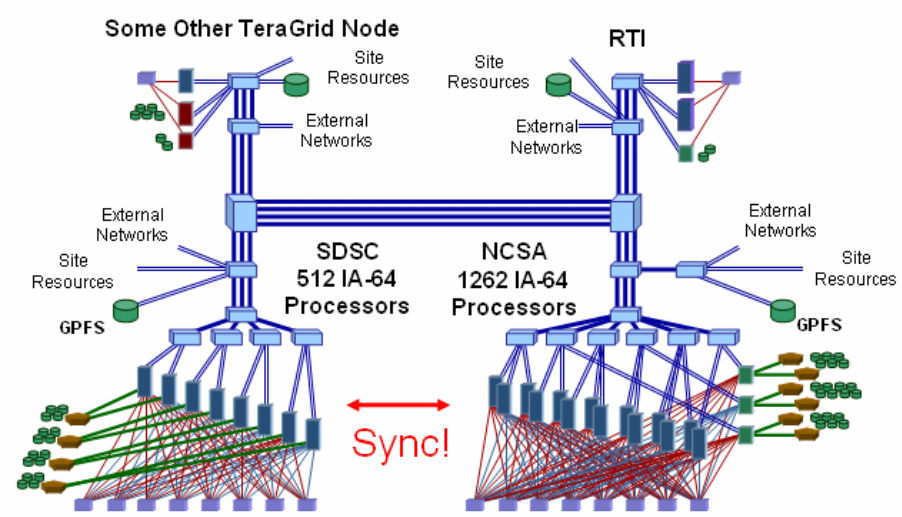

Figure 2. TeraGrid Hardware used in the VBI/RTI MRAC

5 Advance Support for TeraGrid Applications (ASTA)

The ASTA Program aims to ensure that the resources of TeraGrid are optimally utilized for important scientific discoveries and technological innovation. It associates one or more TeraGrid staff members at established minimum time with a user's project for a sustained effort aimed at maximizing the effectiveness of application software and TeraGrid resources to support a project's scientific goals. Multi-clustering; parameter sweeps; ensemble runs; different but linked models on possibly heterogeneous systems; large scale computation with remote data storage or visualization; check-pointing and steering with code and data migration; computation with live data processing and assimilation; and distributed data mining leveraging a central data collection are examples of innovative application scenarios we believe can benefit from the ASTA Project.

6 Use of MPICH-G2

MPICH-G2 is a "grid-aware" implementation of MPI. When a distributed job is started with MPICH-G2, it runs on all clusters in the currently-defined GLOBUS configuration. With MPICH-G2, it is possible to pass a message from any compute node on any cluster to any other compute node in the running GLOBUS cluster configuration. The implementation gives the programmer the illusion of operating in one MPI_COMM_World communicator, i.e. like "normal" MPI running on a single cluster. MPICH-G2 provides tools for identifying and characterizing each cluster in the GLOBUS configuration

7 Use of GPFS

The IBM General Parallel File System (GPFS) is a high-performance shared-disk file system that can provide fast, reliable data access from all nodes in a homogenous or heterogeneous cluster of IBM UNIX® servers running either the AIX 5L or the Linux operating system. It allows all TeraGrid clusters to access the same global file system. The TeraGrid GPFS physical file system is located at the San Diego Super Computer Center (SDSC). Part of the project's experimental design will be to identify GPFS performance issues.

8 Current Status

An initial design has been developed. EpiSims code has been staged and compiled at SDSC and the National Center for Supercomputing Applications (NCSA). Discussions have been held with the MPICH-G2 originator who is also one of the authors of the GLOBUS toolkit. The project has been granted Advanced Support for TeraGrid Applications (ASTA) status, and a TeraGrid support team has been assigned to assist us.

References

[1] Depinay JM, Mbogo CM, Killeen G, Knols B, Beier J, Carlson J, Dushoff J, Billingsley P, Mwambi H, Githure J, Toure AM and McKenzie FE (2004). A simulation model of African Anopheles ecology and population dynamics for the analysis of malaria transmission. *Malaria Journal* 2004, 3:29.

[2] Gu W, Killeen GF, Mbogo CM, Regens JL, Githure JI, Beier JC (2003). An individual-based model of Plasmodium falciparum malaria transmission on the coast of Kenya. *Trans R Soc Trop Med Hyg*. 2003 Jan-Feb;97(1):43-50.

[3] Halloran ME, Piedra PA, Longini IM Jr, Gaglani MJ, Schmotzer B, Fewlass C, Herschler GB, Glezen WP. (2007). Efficacy of trivalent, cold-adapted, influenza virus vaccine against influenza A (Fujian), a drift variant, during 2003-2004. *Vaccine*. 2007 Mar 12.

[4] Handel A., Longini I. M. Jr. & Antia R. (2007). What is the best control strategy for multiple infectious disease outbreaks? *Proceedings of the Royal Society B* 274, 833-837.

[5] Lipsitch M., Cohen T., Murray M., Levin B.R. (2007). Antiviral resistance and the control of pandemic influenza. *PLoS Med.* 2007 Jan 23;4(1):e15 (epub).

[6] Longini I. M., Halloran M. E., Nizam A., Yang Y., Xu S., Burke D. S., Cummings D. A. T. & Epstein J. M. (2007). Containing a large bioterrorist smallpox attack: a computer simulation approach. *International Journal of Infectious Diseases*, Volume 11, Issue 2, Pages 98-108.

[7] Peterson AT, Martínez-Campos C, Nakazawa Y, Martínez-Meyer E. (2005). Time-specific ecological niche modeling predicts spatial dynamics of vector insects and human dengue cases. *Trans R Soc Trop Med Hyg.* 2005 Sep; 99(9):647-55.

[8] Peterson, A. T., and J. J. Shaw. (2003). *Lutzomyia* vectors for cutaneous leishmaniasis in southern Brazil: Ecological niche models, predicted geographic distributions, and climate change effects. *International Journal of Parasitology* 33:919-931.