

Supercomputing in Clinical Practice

by Stefan Zasada, C.V. Gale, Steven Manos and Peter Coveney

We introduce the GENIUS project, part of the revolution in on-demand access to large-scale computing resources for cerebrovascular modelling, surgical planning and intervention.

Traditional medical practice requires physicians to use their judgement and experience to decide on the course of treatment best suited to an individual patient's condition. While training and practice help physicians to hone the ability to choose the most effective treatment for a particular ailment, this decision-making process often fails to take into account all of the data potentially available. Indeed in many cases, the sheer volume or nature of the data available is too much for a human to analyse, and valuable information is discarded from the decision-making process.

The GENIUS project (Grid-Enabled Neurosurgical Imaging Using Simulation) is funded by the UK EPSRC and coordinated by University College London, and involves the Centre for Computational Science, University College London and neurosurgeons at the National Hospital for Neurology and Neurosurgery (Queen's Square, London). The aim of the project is to augment the clinical decision-making process through the development of patient-specific models based on image data, to simulate blood flow within a given patient's brain. These simulations are used in the planning of surgical intervention for arteriovenous malformations and aneurysms. Crucially, this project makes use of large-scale computing resources through on-demand access – a key requirement if patient-specific simulation is to be of use in clinical practice.

Briefly, this project aims to achieve the following scenario. A three-dimensional rotational angiography scan of a patient's brain vasculature is taken in London; a blood-flow simulation is then performed with resources made available on the US TeraGrid, or in the UK (HeCTOR, HPCx or NGS); results in the form of blood-flow visualization are returned in real time to a workstation in the operating theatre control room (see Figure 1) and are either analysed on the fly for urgent surgical intervention, or later as required, for

prognostic and/or surgical planning purposes (see Figure 2).

In practical terms, the patient-specific medical simulation scenario described above requires access to both appropriate patient data and to the infrastructure on which to perform potentially very large numbers of complex and demanding simulations. Resource providers must furnish access to a wide range of resources, typically made available through a computational Grid, and must institute policies that enable patient-specific simulations to be run on those resources. Unfortunately, the traditional 'batch queue' model used by the majority of high-performance computing (HPC) resource providers does not meet the medical computing requirements of this project and others like it. Simulations used to support clinical procedures must be run when required by the clinician: surgical procedures cannot be postponed until a simulation reaches the top of a batch queue. Emergency medical simulations must therefore either pre-empt all other jobs running on a machine, or be scheduled into the clinical workflow through advance reservation.



Figure 1: Cerebrovascular blood flow – von Mises stress field. Blue indicates a stress of 0 Pa and red a stress of 6.7 Pa.

Consequently, the GENIUS project is trialling technology such as HARC and SPRUCE (see links below) that allows batch-processing HPC resources to be used in a more interactive manner. HARC allows users to co-reserve time across multiple resources, which is essential not only when multi-site Message Passing Interface (MPI) jobs need to be run (eg using MPIg; see links below), but also to the scheduling of simulations into the clinical timetable. SPRUCE in contrast allows a user to pre-empt jobs that are already running or are queued on a machine, when it becomes necessary to run an 'emergency' job. Traditionally, such 'urgent' access to computational resources has only been made available in the US TeraGrid for weather-related simulations (eg in the event of a hurricane).

Apart from the requirement for resource providers to support interactive and emergency jobs on their machines, the ability to routinely run computational simulations in support of medical treatment procedures will also require those funding bodies involved in supporting medical research to engage with the HPC community, and develop policies on the provision of HPC access to underpin this growing field of research. Key issues include the impact the requirements of such an infrastructure have on the policies put in place by Grid resource providers, and why resource providers must listen to the requirements of their current user communities and formulate policies that meet the needs of those communities, while also tracking the needs of emerging user groups.

Computational biomedicine, the field within which the GENIUS project operates, is gaining momentum through projects such as the EU FP6 e-Health ViroLab project and the FP7 Virtual Physiological Human Initiative. The latter receives EU funding of €72 million and involves numerous universities, hospitals, organizations and industrial partners throughout Europe. To date however, efforts in this realm have been hampered by a lack of 'joined-up'

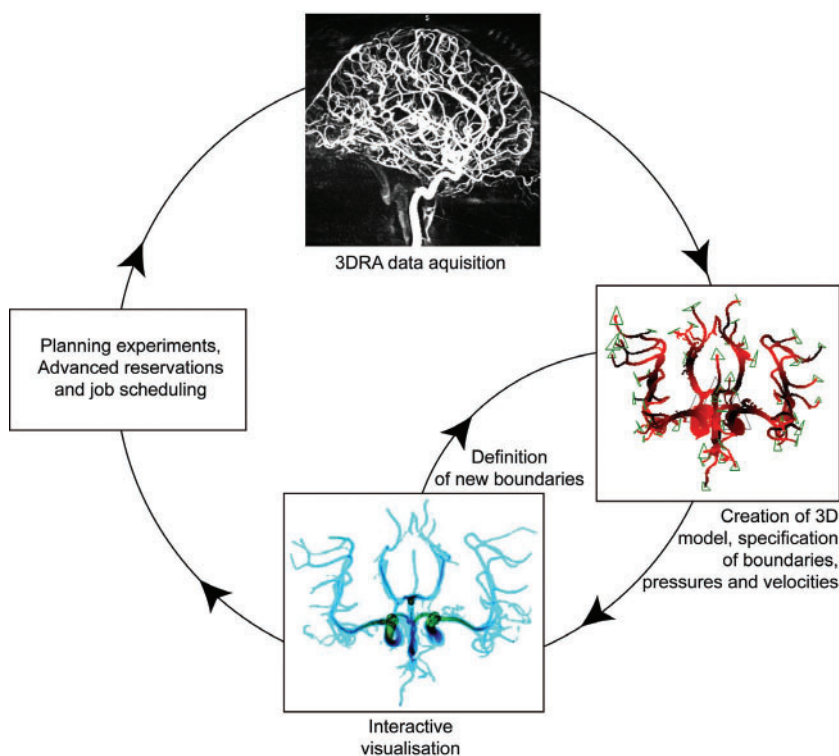


Figure 2: Workflow diagram of GENIUS clinical 'scenario'.

resource provision strategies, at national and international levels, to allocate appropriately scaled resources to projects when they are funded. Such projects often require access to a wide range

of resources as part of a scientific workflow, for example the high-end machines provided by DEISA (Distributed European Infrastructure for Supercomputing Applications) or the TeraGrid, along

with smaller clusters provided by the UK National Grid Service.

The GENIUS project will be complete in December 2009, at which time a unique test case for the use of patient-specific simulation and high-performance computing resources to plan surgical intervention will be available. As described, however, many challenges remain to be resolved before supercomputing at work becomes the premise of the medical doctor.

Links:

- GENIUS project Wiki: <http://wiki.realitygrid.org/wiki/GENIUS>
- VPH Initiative and VPH NoE: http://www.biomedtown.org/biomed_town/vphn
- MPICH-G2: <http://www3.niu.edu/mpi/>
- SPRUCE: <http://spruce.teragrid.org>
- Urgent Computing, CTWatch: <http://www.ctwatch.org/quarterly/archives/march-2008>

Please contact:

Peter Coveney
 Centre for Computational Science,
 Department of Chemistry, University
 College London, UK
 Tel: +44 20 7679 4560
 E-mail: p.v.coveney@ucl.ac.uk

Large-Scale Immune Models and Visualization

by Dimitri Perrin and John Burns

Large-scale computing architectures allow detailed modelling of complex systems. Here we present applications in the field of computational biology.

In silico simulation methods – simulations within computer software – have become indispensable tools in the development of expensive new technology, from aircraft manufacture to nuclear power station development. For various reasons however, modelling and simulation techniques have only very slowly been adopted by the biological research community.

Biological systems are typically complex and adaptive. Given the dynamic nature of these phenomena, it is non-trivial to provide a comprehensive description of such systems and, in particular, to define the importance and

contribution of low-level unsupervised interactions to the overall evolution process. An agent-based approach is presented here.

The immune response, in particular, consists of relatively simple cell-level interactions, but is capable of highly diverse and individual behaviour, with outcomes that are self-evidently more than the sum of their parts. In what follows, a one-to-one reciprocity between autonomous agents and immune cells is proposed. The emergence of tissue-level and body-level patterns is clearly dependent on the relative sizes of cell populations and the balance these

achieve. This is demonstrated, for example, by the course of HIV progression, where the whole immune system collapses once immune cell counts decrease below critical levels.

High temporal granularity is necessary to realistically account for cell mobility and interactions. With current computing resources, this prevents models accounting for every immune cell of every type over the whole body. A compromise between agent diversity and agent population size is therefore required. We have implemented a flexible model, which currently includes viral agents and an agent type for each