
Volume 5 / Issue 1 / 2010 - Cover

Grid Computing and E-Health

Author

Tosh Sheshabalaya,

Grid computing is emerging as a promising solution to some of the most vexing challenges facing e-Health. It also offers a powerful tool for other areas of healthcare such as drug discovery, as well as economic and weather forecasting, earthquake analysis, etc. In all these domains, grid computing easily outclasses traditional IT systems, in terms of its ability to cope cost effectively with their massive demands on computer-processing power, and the intensity of real-time data throughputs.

Grid Computing and e-Health

Whether one speaks of business travellers and tourists outside their home countries, the growing number of elderly with chronic diseases who cannot easily visit hospitals, or urgent, complex cases requiring consultations with a range of specialists, e-health means real-time, secure acquisition and access to extremely vast volumes of data from anywhere, at anytime.

From a pacemaker to an electronic health record, grid computing is seen as one of the ways to address these, as well as a wide array of associated challenges.

Both the robustness and the in-built fault tolerance of grids juxtaposes directly with the demand for 'always live' healthcare applications.

Other than access to distributed databases, the rapid data mining capabilities of grids are seen as enabling tools for a variety of epidemiological and biomedical applications. Grid computing is also proving itself in bioinformatics research, and has several proponents who swear about its advantages in terms of new drug design/discovery as well as personal medicine (or i-health).

At the moment, there are several e-health grid computing initiatives underway in Europe, at different stages – from research through pilot projects to implementation.

Grid Computing in Europe

The European Union has been heavily involved with grid computing right from the start. However, two foundational projects, both funded by the EU's RTD Framework Programme, set the stage for the emerging pan-European grid computing infrastructure, including its e-health dimensions.

These projects are known by their acronyms: BEinGRID and Enabling Grids for E-science.

BEinGRID (Business Experiments in Grid) was launched in June 2006, and concluded at the end of 2009. Its aim was to establish effective routes to foster both adoption of grid computing and stimulate further research into grid computing-based business models.

- One of the key e-health applications of BEinGRID is the RadiotherapyGrid (discussed below).

The Enabling Grids for E-science project (EGEE), which includes sites in the US and Asia, is a more application-oriented successor to a previous project called European DataGrid (EDG). It is considered by many to be the world's largest computing grid. EGEE was developed to support the CERN Large Hadron Collider, which requires storage rates of several gigabytes per second.

- One of the key e-health applications of EGEE is the Health-e-Child project. Another is VPH. Both are discussed below.

E-Health Grid Projects in Europe: A Sample

Radiotherapy Grid

The usual recourse to treatment for the roughly 3 million Europeans diagnosed each year with cancer is external radiotherapy. This uses a Linear Accelerator to attack cancerous tissue with radiation, delivered from several different directions. However, calculating the process for delivering the prescribed doses can be an arduous task: too much radiotherapy is as risky as too little. In addition, treatment also involves calibrating the direction, size and length of radiation dosages, in order to avoid damaging healthy tissue.

The RadiotherapyGrid, spun out of the BEin-GRID foundational project mentioned above, uses grid computing technology to plan the most effective/ personalised treatment for each patient. In complex cases, the Grid has already proved its effectiveness in improving the quality of treatment and helping hospitals save money – and thereby treat more patients effectively.

RadiotherapyGrid, which produces results much quicker than hospitals have been used to, has two core objectives: exploiting massive computational for the verification of radiotherapy plans using accurate computing techniques, and searching for optimal treatments. RadiotherapyGrid is also future proofed. It is designed to be easily extended via further algorithms and tools.

The BEinGRID project was meant to develop business models for the use of grid computing.

The process of contracting with Radiotherapy- Grid is transparent, and based on an automated Service Level Agreement (SLA) negotiation. RadiotherapyGrid is planned to be primarily marketed as a Software-as-a-Service (SaaS) platform at the nearly 1,000 institutions in Europe which offer radiotherapy.

Health-e-Child

Health-related projects based on the above-mentioned EGEE (Enabling Grids for E-sciencE) include a paediatric application called Health-e-Child. This is integrating innovative predictive disease models, complex data visualization and knowledge discovery applications. The ultimate goal is to support clinical decision making in cardiology, rheumatology and neuro-oncology, and develop personalised healthcare for children using database-guided biomedical support. The system has undergone security assessment, and released officially at the end of 2008.

It is currently providing secure and anonymous hosting of health records. Four leading European paediatric hospitals are participating in Health-e-Child: Great Ormond Street Hospital (London), NECKER Enfants Malades (Paris), Istituto Giannina Gaslini (Genova) and Ospedale Pediatrico Bambino Gesù (Rome).

NeurIST

This project involves developing an end-to-end grid IT infrastructure for the management of all processes linked to research, diagnosis and treatment for multi-factorial diseases. Based on a service oriented architecture (SoA), it is targeted at use by medical researchers, clinicians and other healthcare specialists, as well as IT solution vendors and medical suppliers.

neurIST's first phase involves one condition: cerebral aneurysm and subarachnoid haemorrhage. However, its core technologies are designed upfront to be replicable for other diseases and disease groups. The infrastructure consists of datawarehouses, computational analysis services and multiscale, multi-modal information systems at distributed sites.

Akogrimo Mobile Grid Framework

Akogrimo is specifying and prototyping a mobile grid computing infrastructure for a Heart Monitoring and Emergency Management Scenario (HMES). Akogrimo supports establishing a full range of applications - from patient monitoring via ECGs with data transfer to the grid via mobile telephones, to complex procedures such as remote emergency handling.

The HMES service, which aims at the early recognition of heart attacks or apoplectic strokes and rapid access to treatment, has established two baseline cases:

- Ó A permanent cardiac monitoring service with rapid alert triggering and a disease-specific response in an emergency; and
- Ó A non-permanent monitoring device activated by patients suspecting a cardiac problem.

Other than mobility, HMES's key facets are:

- Ó Cross-disciplinary service involving ambulances/ paramedics, hospitals and physicians; as well as network operators and application services providers
- Ó Decision support services to manage high loads of live and historic patient data.

Virtual Physiological Human (VPH)

This project involves a computer simulation model for investigation of the human body as a single complex system – from the cellular level, through individual organs to the whole individual organism. Its eventual goal is to integrate biomedical research across disciplines, and pave the way for personalised medicine (i-health).

VPH necessitates both a massive amount of computing power and demands on data storage and management. Grid computing, including the EGEE foundational project, is being harnessed to meet such needs.

Some key VPH projects, with extremely heavy data and computing processor power demands, include:

ARCH - Patient specific image-based computational modelling for improvement of vascular access in patients on haemodialysis therapy.

ARTreat – Patient-specific artery and atherogenesis model for outcome prediction, decision support treatment, and virtual training.

ContraCancrum – Composite multi-level platform for simulating malignant tumour development and tumour and normal tissue response to treatment modalities and schedules.

EuHeart – Patient-specific cardiovascular modelling and simulation for medical device evaluation and optimisation.

HAMAM – Highly-accurate breast cancer diagnosis through integration of biological knowledge, novel imaging modalities, and modelling.

IMPPACT – Image-based multi-scale physiological planning for radio frequency ablation cancer treatment.

NeoMARK – Advancing current models and methods predict neoplastic reoccurrences, and to apply it to the study of oral cancer. Osteoporotic Virtual Physiological Human – Diagnostic image- based modelling to predict the strength of patients' bones and how this strength will change over time.

PASSPORT – Patient-specific liver modelling combining anatomical, mechanical, appearance and biological preoperative modelled information.

PredictAD – New biomarkers and clinically useful tools for early Alzheimer's disease diagnosis.

RADICAL – Roadmap for enhancing security to protect medical and genetic data.

VPH2 – Patient-specific computational modelling and simulation of the human heart to assist the cardiologist and cardiac surgeon in defining the severity and extent of Left Ventricular Dysfunction (LVD), with or without ischemic mitral regurgitation.

Technical Challenges for Grids in the Healthcare Context

While its operational/functional features are driven by technology (see box), grid computing in an e-health context has to meet several standards on security. These impact upon any healthcare information technology system, but have even more stringent demands in terms of grid computing.

Some of the key challenges are :

Access: Different users of a grid (hospitals, GPs, specialists, authorities and patients) will require flexible controls on access.

Authentication: Every healthcare IT system faces the problem of verifying the source and authenticity of requests for information. Distributed grid computing, by its very nature (anytime, anywhere) makes such a task even more demanding.

And yet, authentication should not be a significant burden on users, especially medical professionals on the move. One potential solution is seen to lie in biometric sign-up systems.

Audit trails: An audit trail is already a regulatory requirement, especially to ensure that (the chain of) those modifying a patient's healthcare data are identifiable and traceable.

Security breach detection: Given the dispersal of both hardware and application resources across a grid computing network, one challenge is IP spoofing. The challenge is to devise a system for both user and server certification, which allows for mutual authentication during any transaction, without producing a drag on the speed of the overall system.

Once again, given the distributed nature of the grid, preventive monitoring of security breaches is a major technical challenge.

Data encryption and integrity: Following mutual authentication by both user and server (see above), the next step in grid computing is devising an encryption key for local data transfers, alongside VPN technology for remote transfers.

Given the sensitive nature of healthcare information, data integrity issues in a grid involve both encryption as well as digital signatures. This is to ensure that data outputted from a system is reliable, not only in terms of reading but also as far as modifications are concerned. Most grids also seek to retain a back-up version of the data, prior to modification.

The Future: Technology Development and Policy

In addition to the issues described above, it is crucial to take account of the fact that neither the challenges of e-health, nor its enabling technologies, are static.

To make e-health a reality, grid computing initiatives have to be closely tied into progress in terms of interoperability of healthcare IT systems and integration with Electronic Health Records (EHRs).

A related challenge will be to grid-enable medical devices, which face their own interoperability and interconnectivity conundrums with regard to healthcare IT applications, in what is bound to be an artificial hardware-software segregation.

SHARE (Supporting and structuring HealthGrid Activities & Research in Europe)

The EU SHARE project, which concluded in 2008, investigated the technology and policy challenges mentioned above, and drew up a roadmap for accelerating deployment of grid computing in the healthcare area.

As expected, key issues raised by SHARE concerned security, data protection and privacy, and the trades-off between easy access and security, not least in terms of conformity with data protection laws. SHARE's work is being continued by the EU's new e-Infrastructure Reflection Group (e-IRG).

One looming challenge for policy makers will be how to formalise patient consent in terms of being fully informed about who will have access to their records across a health grid, and how this would be used.

The inherent conflict - is making anonymous the vast volumes of data in a grid, while permitting rapid identification of its origin for emergency treatment and personal medicine. Such a task is not going to be technologically straightforward.

Grid Computing: The Essentials

The System

In technical terms, grid computing consists of large-scale, distributed cluster computing, coupled to network- distributed parallel processing. The network interface is a traditional one such as an Ethernet. This sets grid computing apart from a traditional parallel-processing supercomputer, which uses many processors connected by a local high-speed computer bus.

Again, unlike the customised hardware of a supercomputer, usually built in a small series, the bulk of components in a grid computer are commodity off-the-shelf products. Last but not least, scalability is inbuilt into the grid because of a lower need for connectivity between nodes as compared to the capacity of the Internet.

The Process

Fundamentally, the grid computing process consists of breaking up and assigning blocks of a specific software program between hundreds or thousands of independent computers, and thus harnessing their additional power – not least when this is idling or functioning at below maximum rating. The architecture of grid computing extends from managing computing processors to data storage and transfer, security, remote monitoring, as well as a toolkit for developing a range of customised services for specific applications.

One key element of the process is specific middleware, to allow sharing heterogeneous resources.

The users of grid computing extend from one large entity at one or more sites to public collaboration across many organisations and countries. Indeed, for cross-border collaborative initiatives like Europe's ehealth program, grid computing seems a natural.

One of the most famous grid computing networks is SETI@home, which used more than three million (mainly personal) computers to search for extra-terrestrial intelligence (ETI).

Its Pioneers

The term 'grid computing' is now just over ten years old.

Conceptually, it was defined in a book 'The Grid: Blueprint for a new computing infrastructure' (Morgan Kaufmann Publishers, 1999) as a means to make computer power available on demand from a variety of sources – rather like an electric power grid. The authors of the book were two of grid computing's three fathers, Ian Foster (Argonne National Laboratory at the University of Chicago), and Carl Kesselmen (Associate Professor of Computer Science at the University of Southern California); and Steve Tuecke (also from Argonne National Laboratory).

These three created what was called the Globus Toolkit. Until the present, the Globus Toolkit remains the de facto standard for designing and building grids. However, one of its core middleware components is facing competition from newer entrants.

Recent Developments

Technology, meanwhile, marches forth. The grid computing market has several sub-segments: middleware, grid-enabled applications, utility computing (provision of grid computing and applications as service), and more recently, software-as-a-service (SaaS).

There are also close parallels between grid computing and cloud computing. Indeed, new systems such as AppLogic from 3tera overlap the two fields. A more compelling case for such a fusion is the fact that some of the pioneers of grid computing are now playing lead roles in cloud computing.

Published on : Mon, 5 Apr 2010