

International Journal of Computational Methods
© World Scientific Publishing Company

GRID SIMULATION SERVICES FOR THE MEDICAL COMMUNITY

GUNTRAM BERTI

GUY LONSDALE

JENS GEORG SCHMIDT

*C&C Research Laboratories, NEC Europe Ltd., Rathausallee 10, D-53757 St. Augustin,
Germany*

{berti|lonsdale|schmidt}@ccrl-nece.de

http://www.ccrl-nece.de/

SIEGFRIED BENKNER

Institute for Software Science, University of Vienna, Austria

sigi@par.univie.ac.at

D. RODNEY HOSE

JOHN W. FENNER

DAVID M. JONES

Medical Physics, University of Sheffield, UK

{d.r.hose|j.w.fenner|d.jones}@sheffield.ac.uk

STUART E. MIDDLETON

IT Innovation Centre, University of Southampton, UK

sem@it-innovation.soton.ac.uk

GERT WOLLNY

MPI for Human Cognitive and Brain Sciences, Leipzig, Germany

gert.wollny@acm.org

Received (Day Month Year)

Revised (Day Month Year)

The first part of this paper presents a selection of medical simulation applications, including image reconstruction, near real-time registration for neuro-surgery, enhanced dose distribution calculation for radio-therapy, inhaled drug delivery prediction, plastic surgery planning and cardio-vascular system simulation. The latter two topics are discussed in some detail. In the second part, we show how such services can be made available to the clinical practitioner using Grid technology. We discuss the developments and experience made during the EU project *GEMSS*, which provides reliable, efficient, secure and lawful medical Grid services.

Keywords: medical simulation; grid computing; bio-medical engineering.

1. Introduction

In the SimBio project [Simbio], we used the expression “Bio-numerics” for numerical simulation and modeling of bio-physics problems related to human body parts of interest for specific medical treatment — as a means of distinguishing the work from bio-informatics, whose focus is generally proteomics or genomics. The key distinguishing factor compared with image processing or image reconstruction in the same applications arena is the use of computational methods for predictive purposes — providing physically accurate (up to the possible modeling accuracy) information not included in medical images themselves. An example which clarifies the difference: by visualization one could investigate whether or not a prosthesis would fit for a given patient, even during movement (with some simple assumptions); with bio-mechanical modeling one could investigate the resulting forces on the prosthesis and surrounding body parts, allowing predictive statements to be made on durability or perhaps even comfort or pain for the patient.

It is fair to say that enormous progress has been made in recent years (aided by the increases in performance of computing platforms) and that numerical modeling is now able to provide realistic (and validated) predictions of very complex phenomena. However, there is of course, as in other areas, a wide “proximity to black box product” spectrum in the Bio-numerics area. There are examples of numerical simulation embedded within current tools that are, or could be, used in a clinical setting now. One example is the radio-surgery / radio-therapy planning (see Sec. 2.1.1) included in the set of applications addressed by the GEMSS project [GEMSS]. Another example, taken from the SimBio project, is electromagnetic source localization — a non-invasive diagnosis tool using an inverse method based on external EEG/MEG measurements. Many areas of development in numerical human modeling are already at the stage that they can be used by medical researchers, (who are not themselves modeling/software developers), as tools for investigation into causes of medical problems (or processes leading to problems) and treatment procedures. We discuss these and other examples of medical simulation in the first part of the paper (Sec. 2), giving a more in-depth treatment of facial reconstructive surgery (Sec. 2.2) and cardio-vascular system modeling (Sec. 2.3).

The question, however, is how to make such applications accessible to the clinical practitioner. This involves on the one hand ergonomic user interfaces, which are discussed in the first part of the paper, and on the other hand, access to high-performance computing (HPC) hardware, which is typically not available at clinics nor easily usable by non-experts. While there are bio-physics applications, (e.g. *in silico* screening or docking for drug design), that are suitable for simple meta-computing approaches (i.e. task-farming in standard hardware across wide-area networks), most of the envisaged bio-numerics applications are more demanding in terms of computing resources. The common aspect reported on here is computational Grid services — software solutions that are developed for or adapted to medical tasks, and that are pre-installed on large computing systems. Bio-numerics

applications often require HPC systems or can benefit from the use of more computationally intensive algorithms or methods that deliver increased accuracy (and in doing so, medical benefit). Grid can be used to provide access to appropriate computational services and deliver these to medical users. The major challenges are to ensure that services can be delivered into the user's workplace in an appropriate, ergonomic manner and that security, policy and legal constraints related to the use of patient data are fulfilled. These issues are discussed the second part of this paper, which deals with necessary Grid technology. Here we report on our experience in the GEMSS project, which was a dedicated effort for developing Grid services for medical simulation, in a way conforming to security standards and (EU) law.

2. Part I: Medical Simulation

The scope of applications for medical simulation is vast. In this first part, we aim to give an idea of the diversity of the field, concentrating on work we have been involved with, and briefly relating it to its grid context. Two areas — facial reconstructive surgery simulation (Sec. 2.2) and cardio-vascular system simulation (Sec. 2.3) — are treated in some depth in order to provide more insight into the typical challenges associated with medical simulation.

2.1. A Quick Tour of Medical Simulation Applications

2.1.1. Radiosurgery planning

The Gamma Knife[®] is a radio-surgery device that uses 201 targeted ⁶⁰Co sources to treat brain lesions. At present, the radiation dose distribution is calculated rapidly using the GammaPlan[®] software which employs a generic geometric model. However, superior results may be obtained using complex, compute-intense Monte Carlo simulations, cf. Fig. 1 which shows that the improvements are especially apparent at the boundary between materials with different attenuation coefficients [Fenner *et al.* 2005]. These simulations are particularly suited to the Grid because of their parallel scalability and their large computational demands.

2.1.2. Inhaled Drug Delivery Simulation

The process of pulmonary drug delivery is affected by many factors, such as inhaler design, formulation of the medication, airway geometry, and the drug absorption process. Computationally demanding 3D fluid dynamics simulation is used to calculate the pattern of air flow and predicted absorption rate (Fig. 2). Computing the absorption of inhaled drugs is of great interest to the pharmaceutical industry. Corresponding parameter studies require substantial computational resources.

4 *G. Berti et al*

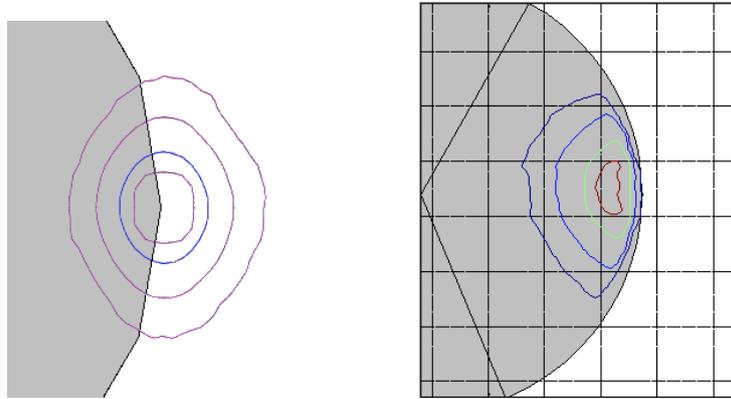


Fig. 1. Comparison between radiation dose distribution calculated using GammaPlan[®] (left) and a Monte-Carlo method (right).

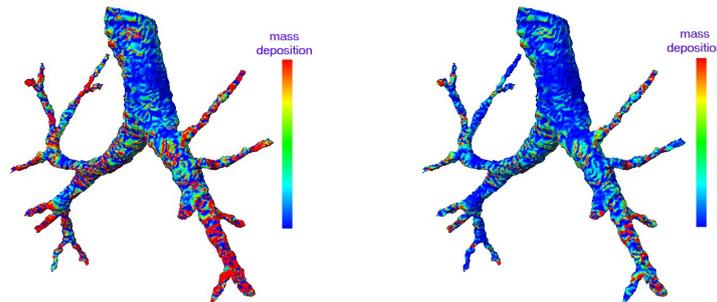


Fig. 2. Deposition of water droplets in the lungs — smaller droplets ($4 \mu m$ - left vs. $8 \mu m$ - right) penetrate deeper into the lungs.

2.1.3. SPECT 3D Image Reconstruction

Tumor diagnosis and monitoring of metabolism are the main tasks of in vivo diagnosis in nuclear medicine by visualization of distribution of radioactive tracer in the human body. Although SPECT (single photon emission computed tomography) reconstruction suffers from low spatial resolution and poor signal-to-noise ratio compared to modern x-ray CT and MRI, it provides complimentary functional information, and is indispensable in modern clinical diagnosis. The diagnostic procedure requires the patient to receive an ionizing radiation dose which provides the radiation necessary for acquisition of multiple projections by the gamma camera. The projection data is computationally reconstructed for subsequent reporting and diagnosis. Standard reconstruction algorithms based on filtered back-projection (FBP) find extensive use in current clinical practice but are only applicable to single slices and deliver reconstructions of limited quality. Modern iterative, three-dimensional

algorithms like ordered subset – maximum likelihood (OSML) [Backfrieder *et al.* 2003] offer benefits (cf. Fig. 3) because they encompass technical and physical constraints of the imaging process, but are computationally expensive and thus a natural candidate for Grid computing.

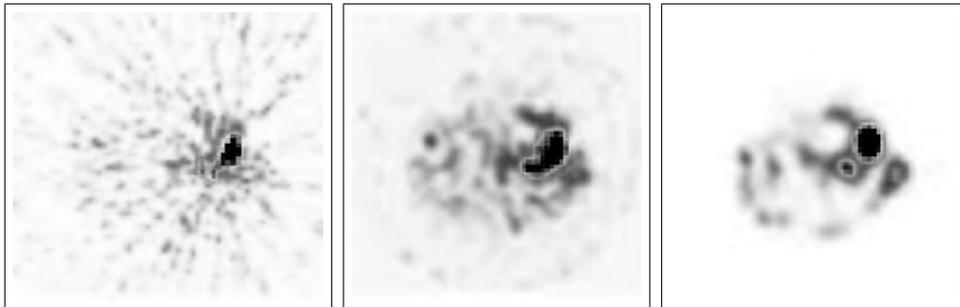


Fig. 3. SPECT image reconstruction comparison (I^{131} study of endocrine system, 60 slices with 64×64 voxels of $8.84mm^3$): Left: 2D filtered back projection (FBP), middle: 2D iterative reconstruction without attenuation correction (IRNC), right: 3D iterative reconstruction using an ordered subset – maximum likelihood (OSML) algorithm used in GEMSS, giving the best results.

2.1.4. Source Localization from EEG/MEG Data

A common clinical task in neurology and neuropsychology is to find realistic electromagnetic source distributions in the human brain. This is used for instance to guide surgical treatment of epilepsy, or for cognitive research. The point of departure are EEG measurements giving electrical potentials on the surface of the head. This results in an inverse problem, the solution of which may require several thousand forward problems to be solved, typically by either the finite element or the boundary element methods [Wolters *et al.* 2002]. Consequently, large computational resources are required.

2.1.5. Knee Mechanics

Understanding and modeling knee joint kinematics is important for testing and optimizing implants, such as artificial menisci, or for analyzing the gait cycle [McCarthy *et al.* 2002]. A corresponding computational model is fairly complex, consisting of 3D, 2D and 1D components (for ligaments). Modeling of knee bending or gait analysis leads to a contact-dominated simulation, see Fig. 4.

2.1.6. Neurosurgery support

In neurosurgery, the resection of brain tumors may be planned on the basis of high resolution pre-operative anatomical magnetic resonance (MR) images of the

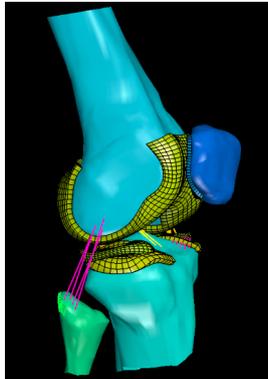


Fig. 4. Computational knee model. Ligaments are modeled as 1D elements.

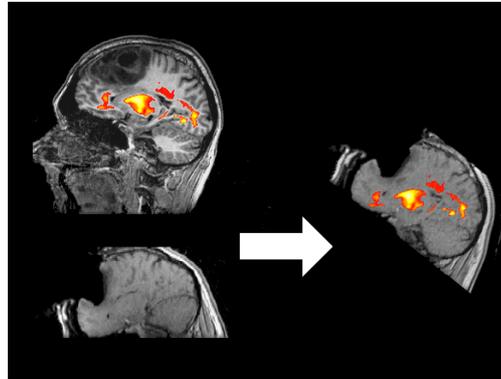


Fig. 5. Neurosurgery support: A pre-operative anatomical and functional MR image (left top) and an intra-operative low-resolution anatomical MR image are registered to transfer the functional information to the inter-operative image data.

patients head. To gain more information about the function of regions adjacent to the tumor, a functional MR image (fMRI) may additionally be acquired. This information is useful in assisting the surgeon's navigation during the intervention and to minimize the risk of potential functional damage.

However, after the opening of the skull several effects take place that lead to non-linear distortions of the brain, collectively called the brain shift phenomenon. Thus, pre-operatively acquired functional information can not easily be mapped onto intra-operatively acquired anatomical images. This is the major shortcoming of image-guided surgical planning based on pre-surgically acquired fMRI data, because the occurrence of surgically induced deformations invalidates positional information about functionally relevant areas.

This problem is addressed by non-linear registration of pre-operative fMRI images to intra-operative MR Images acquired by an Open-MR scanner (cf. Fig. 5). So whenever an intra-operative dataset is acquired, the following image processing chain must be executed [Lippmann and Kruggel 2005]:

- (1) transfer of anatomical images from the scanner and conversion into a machine-independent data format
- (2) correction of intensity non-uniformities in the scan data
- (3) linear registration of the (low-resolution) intra-operative scan with a (high-resolution) pre-operative scan
- (4) intensity adjustment between both scans
- (5) non-linear registration of both scans to yield a 3D deformation field
- (6) application of the deformation field to the pre-operative fMRI datasets
- (7) overlay of the deformed functional information onto the intra-operative data

- (8) conversion and transfer to the presentation device (e.g., monitor, surgical microscope)

Due to the requirement of processing the computational part of the chain (excluding image acquisition) in approximately 10 minutes, large computational power is needed, which is best delivered by a Grid service. A particularly challenging requirement for a quality-of-service guarantee (see Sec. 3.5.2) is the need to have immediate availability of this service, as it occurs in the middle of the surgical intervention.

2.2. Detailed Example 1: Maxillo-facial Surgery Planning

2.2.1. Medical Background

Cleft lip and palate are among the most frequent inborn malformations. A resulting maxillary hypoplasia can be treated by *distraction osteogenesis*: During an operation the appropriate bony part of the midface is separated from the rest of the skull (*osteotomy*) and slowly moved into the ideal position by way of a distraction device (cf. Fig. 6). Thus even large displacements over 20 mm can be treated effectively.

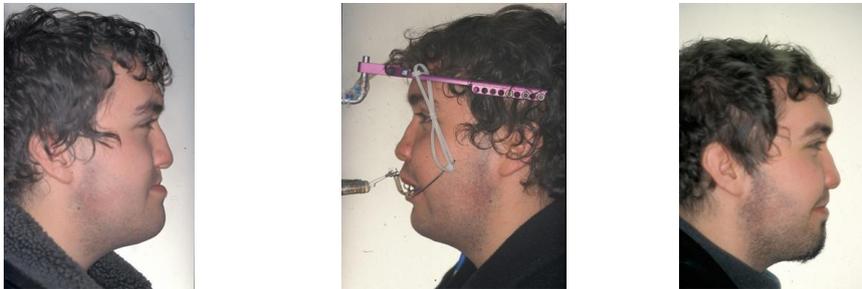


Fig. 6. Patient before treatment (left), in the middle of treatment with distraction device mounted (middle), and some time after the treatment. The cable connects to a data logger device, which records the pulling forces every 6 seconds. It is planned to use this information for validating the nonlinear simulations.

A critical point in this procedure is osteotomy design and predictions of the resulting outcome with respect to aesthetics. Another important question is the optimal adjustment of distraction forces, as on one hand, it is desirable to pull as fast as possible, while on the other hand, excessive stresses cause unnecessary pain for the patient.

In current clinical practice, planning is solely based on CT scans and the surgeon's experience. However, given a sufficiently detailed model of the patient's head, it is possible to predict the outcome of such surgery by using computational structural mechanics. This involves a toolchain which helps a surgeon to use numerical

simulation to plan his treatment. Starting from high-resolution CT images, the different materials (bone and soft tissue) are identified (*segmentation*), a surface mesh is generated and used by the surgeon to specify bone cuts and displacements, which are further processed to produce a valid 3D FEM model. This model is fed into a remote simulation application, and the result is presented to the surgeon (Fig. 7).

A number of researchers have tackled this challenging problem before. In Koch [2001], a system based on linear elasticity is described, where osteotomies are specified on 2D pictures. Schutyser *et al.* [1999] use a 3D “virtual knife” and emphasize the real-time aspects of simulation, trading accuracy for speed. Zachow *et al.* [2001] use a specialized version of the Amira visualization system for most parts of the toolchain, including mesh generation. Simulation is restricted to linear FEM models. In Gladilin [2003], these linear models are extended to first non-linear FEM simulations of the distraction process, using a St.Venant-Kirchhoff material model. Chabanas *et al.* [2004] compare linear and non-linear simulation of maxillo-facial surgery; however, the small displacements used do not permit a conclusive judgment.

Our approach differs from previous efforts by focusing on autonomous usage by non-technical users. This goal entails the offering of transparent access to high-performance platforms, eliminating hardware-related constraints. The corresponding Grid solution developed in GEMSS opens the door for compute-intensive, high-fidelity numerical methods, which we consider indispensable when striving for accurate, reliable solutions. In particular, to cope with the large displacements present in the model, we decided to use non-linear approaches for the FEM simulation.

A major challenge in this context is the *automatic* generation of *valid* physical models from imperfect geometry descriptions such as segmented medical images, which moreover are often polluted by imaging noise such as metal streak artifacts. Furthermore, non-linear approximations are generally rather sensitive to model imperfections, aggravating this problem. More details on difficulties arising in this context and some corresponding solution approaches are described below.

2.2.2. *Outline of the Maxillo-facial Surgery Planning Toolchain*

A complete toolchain for the simulation of maxillo-facial surgery encompasses a considerable number of individual steps [Schmidt *et al.* 2004]. First of all, we let the surgeon specify *landmarks* on the CT dataset, using a specialized tool displaying iso-surfaces at per-landmark selectable gray-values, on which landmarks can be picked. The resulting model-specific expert knowledge is integrated into the subsequent processing and has proved immensely useful.

Next, the raw CT image needs to be segmented, yielding a labeling of the non-background voxels into material classes, encompassing at least bone and soft tissue. Segmentation is a very difficult process for which no general automatic procedures exist. A simple method is based on thresholding, which is rather successful on our

case if only bone and soft tissue classes are required. Thresholding fails when more detailed classification is needed, for instance muscles and nose cartilage, because the corresponding gray value ranges tend to overlap. We use landmark information to enforce separation of upper and lower jaws at this stage.

Further problems for segmentation are imaging artifacts stemming from metal teeth implants and braces, which lead to streak-like spikes in the image. We use morphological image filters to mitigate the effect of these artifacts, which could otherwise severely distort the simulation.

Following segmentation, a mesh of the bone surface is generated. Using this mesh, the surgeon specifies bone cuts and displacements of bone components in a virtual osteotomy tool. The outcome of this virtual osteotomy is used to produce a complete FEM model of the surgery problem. In order to obtain a valid model, we must ensure that the cuts indeed separate the specified bone components from the rest of the skull. We use specialized morphological image processing, guided by landmarks, to enforce separation on the image level.

On top of the accordingly manipulated image, a volumetric 3D mesh of the whole head is generated. Our mesh generation software can produce 3D meshes directly from segmented images, without relying on intermediate 2D surface meshes [Berti 2004]. Either forces or prescribed displacement boundary conditions specified by the surgeon are applied to the model, complemented by (somewhat arbitrary) zero displacement boundary conditions which are automatically derived from appropriate skull landmarks. A typical mesh contains some 400,000 elements (hexahedral elements with 2mm resolution), but may be several times larger. Prior to volume mesh generation, we use landmark information to segment (and exclude) brain, which has practically no influence on the simulation result.

Using the GEMSS system, the computational model is transferred securely to a remote HPC server. Here, it is first partitioned into an appropriate number of processes according to model size and the agreed quality-of-service requirements (see Sec. 3.5.2), and then passed to our parallel structural mechanics simulation. After the simulation has finished, GEMSS middleware transfers results back to the client where they are visualized.

2.2.3. *Finite Element Modeling*

The nature of the problem calls for non-linear modeling in several respects. First, because of the large deformations in the order of about 20mm, the strains cannot be approximated linearly from the deformations, resulting in a geometrically non-linear deformation. Second, as the strains may be large, we use a non-linear material formulation, viz. a Neo-Hookean material law. Third, in the case of living material, the transient adaptation of the material to the loads is of visco-elastic type. When using force-based boundary conditions, these properties cannot be neglected. Only with visco-elastic material laws is it possible to model the transient process correctly, accommodating the overcorrection performed by the surgeon and the subsequent

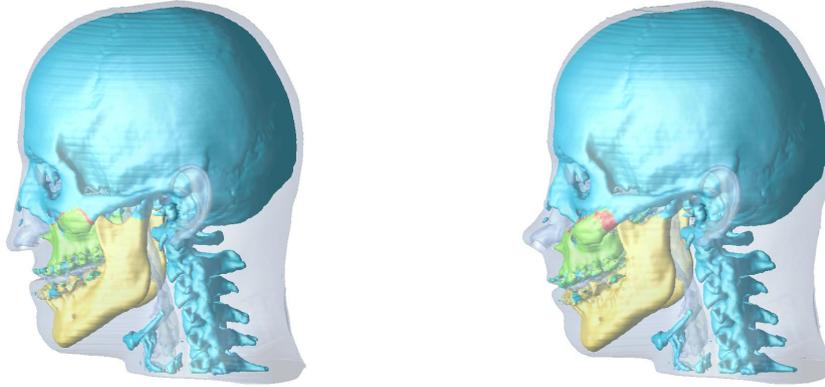


Fig. 7. Patient before treatment (left) and simulated surgery (right), using volume rendering of original and deformed CT image, together with a coloring of the main bone components. Note the spikes around the mouth resulting from image artefacts.

relaxation of the material into the final position when the distractor is removed, i.e. when the forces cease to be applied. A purely elastic material law would predict a return to the initial position, which is evidently wrong.

Currently, we work mostly with tri-linear hexahedral elements, but quadratic hexahedral or tetrahedral elements are also possible. In our experience, hexahedral elements show the best stability in the presence of mesh distortions.

Solving the linear systems arising in the non-linear solution steps is problematic. We employ a GMRES Krylow solver preconditioned by the algebraic multigrid (AMG) method of HyPre [HyPre]. The solution accuracy obtained suffices for the convergence of the Newton method used for the solution of the non-linear systems, yet convergence of the linear solver stops at a certain point, which is a well-known problem for the application of AMG to systems arising from structural mechanics.

We noted that the stability of the Newton method critically depends on the absence of certain mesh anomalies which may occur when using a discrete and approximate geometry description like that provided by segmented medical images. A particular problem is large jumps in material properties in combination with non-manifold connection of elements of “hard” material. An example is a small group of bone elements embedded in soft tissue that is connected to the main bone component only via a vertex but not via a face. Such situations have to be detected and removed.

On a more global scale, the reliability of the simulation depends on a topologically correct model reconstruction. For instance, segmented medical images typically exhibit missing separation of physically disconnected parts, like upper and lower teeth, cheeks around the teeth or upper and lower lips. While we have had some success in automatically ensuring separation of some of these anatomical structures,

the topic remains a challenge.

2.2.4. Future Work

Our research will proceed in three main directions: First, strengthening the *robustness* of the toolchain. This includes topologically correct segmentation, mesh optimization for non-linear simulation, possibly including remeshing, and improvement of the convergence of linear and non-linear solvers.

Second, improving the *modeling*, through better capturing of the nose cartilage and segmentation and modeling of anisotropic tissues like skin and muscles. It might also be possible to account for the *growth* of living tissues like bones and the various kinds of soft tissue.

Third, a thorough validation using real surgical outcomes and measured forces is required. Currently, only a simple validation has been undertaken, based on a few landmarks on the nose, which is insufficiently modeled at the moment. The greatest problem with validation is the lack of unambiguous landmarks in the interesting regions like the cheeks where large deformations occur.

A more advanced opportunity for further research lies in the modeling of *functional* behavior after a surgery, in particular, addressing question like ‘How well a patient will be able to speak, breathe, chew or smile after a treatment?’. This, however, will demand substantial research effort.

2.3. Detailed Example 2: Cardio-vascular Simulation

2.3.1. Application Background

The GEMSS cardiovascular simulation environment offers a set of tools that can be used to evaluate local or global haemodynamic characteristics. This is achieved by coupling a one-dimensional representation of the systemic circulation, in which global pressure and flow characteristics are computed, to a local three dimensional computational fluid dynamics (CFD) representation of the region of interest for detailed haemodynamic evaluation. The 3D computational fluid dynamics analysis requires a series of processing steps:

(1) *Definition of the geometry and mesh generation*

The local geometry is likely to be defined by segmentation of a medical scan to identify the boundaries of the vessel. The degree of difficulty of the segmentation problem depends on the imaging modality. There are many tools available for segmentation and these might or might not require high performance computing. If the object of the study is to evaluate the performance of a medical device, such as a vascular graft or stent, then the geometry of the device might be defined in a CAD system, and in this case the device geometry will need to be integrated with the scan geometry. An overview of the process of generation of a CFD mesh from a medical scan is presented by Steinman [2002]. For the

purposes of GEMSS it is assumed that segmentation is carried out locally using the tools developed in the SimBio project [Barber and Hose 2005]. In the latter approach, a ‘reference patient’ data set is created that is easily adapted to an entire class of similar candidates. The mesh geometry of the reference is morphed to any individual using data obtained from image registration. The process begins with an idealized representation of the real geometry, which is defined in terms of geometric primitives, such as cylinders, and then meshed. From the mesh, a pseudo-binary image is created by identifying which points on a regular 3D grid lie inside the mesh surface (intensity 1), and which lie outside (intensity 0). The scan data for the reference patient is then segmented using any appropriate technique to produce a second binary image. Because this need only be performed once, the segmentation is performed manually. The two binary images are blurred using a Gaussian filter and then registered to obtain a mapping from the idealized geometry to the real geometry. By applying the registration mapping to the idealized geometry, a mesh is obtained for the reference patient. This technique generates a clinically realistic vascular geometry of appropriate mesh density in which the inlets and outlets can be automatically identified by reference to the original geometric primitives.

(2) *Inlet and outlet boundary conditions*

For many cardiovascular analyzes the purpose of the study will be to determine the effect of a particular intervention (whether surgical or by the introduction of a medical device) on haemodynamic characteristics. The problem is that the classical CFD approach of specifying some combination of flow and pressure at the inlet and outlet boundaries is generally inappropriate. The local element of the circulation under investigation is part of a system, and intervention in one area has global consequences. The characterization of the association between the limited set of non-invasive pressure and flow measurements that might be available from routine clinical examination, and thus from the electronic health record (HER), and global arterial function remains an area of significant research interest (e.g. Segers *et al.* [1997]; Segers and Verdonck [2002]; Stergiopoulos *et al.* [1998, 1999]; Kass [2002]). There is obvious benefit in combining a local three dimensional model with systemic models that can supply approximations for the boundary conditions of the 3D domain, particularly in the absence of measured local data. Coupling of a 3D model to a lumped parameter model (a Windkessel model) was adopted in the BloodSim project and in the COPHIT project, and subsequently applied to the analysis of prosthetic heart valve mechanics [Hose *et al.* 2005]. A more sophisticated coupling of a local 3D model to a systemic model — either lumped parameter or 1D — is described by Quarteroni and Venezian [2001] and Formaggia *et al.* [2001]. Numerical solution of the problem involves iteration, possibly including a relaxation step, between solution of the 3D and 1D sections until a global solution is obtained at each time step. The authors indicate that spurious reflections that occur at the bound-

aries of the 3D model can be greatly reduced using this approach. Lagana *et al.* [2005] describe an application in which a detailed 3D model of the aortic arch and connecting vessels is coupled to compartments representing the remainder of the systemic circulation. Migliavacca *et al.* [2004] describe an application in pediatric cardiac surgery. In GEMSS the simple assumption that velocity and pressure distributions at the boundaries of the three dimensional domain are those associated with the Womersley solutions for developed oscillatory flow in a tube was made, and the domain coupled to one dimensional representations of elements of the systemic circulation. With careful choice of the integration procedure in the one dimensional elements, this approach proved successful: benchmark simulations exhibited no unexpected reflections at the interface between 1D and 3D elements and continuity of wave propagation through the domains was observed. For demonstration purposes the circuit model described by Westerhof *et al.* [1969], featuring 121 arterial segments represented as combinations of passive electrical circuit elements, has been implemented as one option in the GEMSS cardiovascular simulation suite: alternative systemic descriptions are possible and selectable in the ‘front end’ options. The system has been designed to be extensible so that new compartment representations can be defined if required. During the test phase it was demonstrated that the GEMSS solver produced the results reported by Westerhof for the physical electrical model that he built and tested. Within the GEMSS system a state variable approach has been taken to the generation and solution of the 1D elements. The entries in the matrices describing the 1D system are pre-computed based on its electrical analogue description. Within a 1D element an exact analytical solution assuming a linear variation of pressure in time at each timestep is adopted.

(3) *Wall conditions*

The traditional approach in CFD analysis of elements of the circulation has been to adopt the no-slip condition with rigid boundaries. This approach remains acceptable for initial evaluation of many systems. More recently several authors have investigated fluid solid interaction (FSI) for application to cardiovascular systems (e.g. [Zhao *et al.* 1998; Hose *et al.* 2003; de Hart *et al.* 2003]). One problem with FSI in this context is that the support conditions for the structural phase of the analysis are rarely known. For problems that are dominated by local dilation this might be acceptable, but for systems such as the aorta or the coronaries, exhibiting large global motions, FSI is difficult to justify. An intermediate approach, in which wall motion is imposed based on dynamic image segmentation, is described by Jones *et al.* [2004]. For the purposes of the GEMSS applications it is assumed that no structural analysis is required.

(4) *CFD analysis*

In the GEMSS project it is assumed that all pre-processing (building of the mesh and prescription of boundary conditions, including the state matrices and

any wall motion) and all post-processing is done locally. The compute intensive step is assumed to be the computation of the transient local haemodynamics (with interaction with the system equations). This step is performed remotely on a GRID platform, using the CFX CFD solver, with the facility to interrogate the current status of the remote job at any time.

(5) *Post-processing*

On completion of the CFD solve the results are returned to the local platform for post-processing.

2.3.2. *Application to idealized femoral artery stenosis*

One of the 3D models used to test the GEMSS environment was an idealized stenosis model developed by one of the GEMSS partners (ASD GmbH). Previously this had been run with simple flow or pressure boundary conditions. We improved the boundary conditions by coupling compartment models to represent the circulation upstream and downstream of the stenosis.

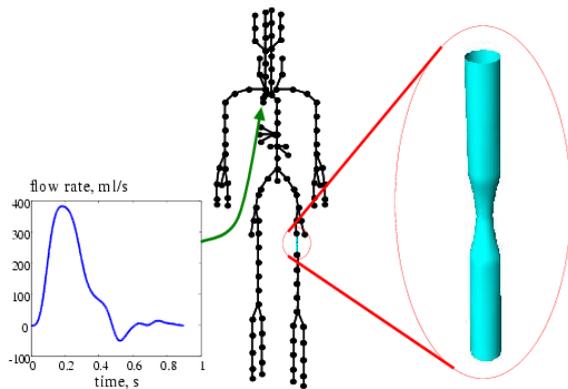


Fig. 8. Stenosis model coupled to compartment models

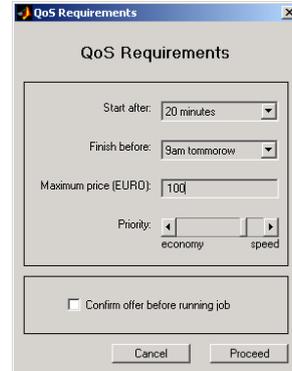


Fig. 9. Specification of Grid Job Requirements

The workflow for this problem is as follows. The 3D geometry of the stenosis was defined and meshed using CFX software, and the inlet and outlet boundaries identified. The 3D model definition is saved to a file. This file is then loaded into the cardiovascular software, which automatically identifies possible boundaries for coupling. The user can choose between a range of possible compartment types for coupling to the 3D model. These include simple resistances, Windkessel models and Westerhof model sections.

For testing, we chose to couple two Westerhof model sections - one to the inlet

and one to the outlet. The section to be coupled is identified using a user interface which displays a graphical representation of the circulatory system. The outlet model is passive, but the inlet model is driven by a flow rate representing the cardiac output. The user can specify this flow rate in a number of ways. In this example a clinically-determined flow rate was used. Finally, relaxation parameters are specified for each of the boundaries. These parameters affect the coupling between the 3D model and the compartments. Smaller values give smoother coupling, but at the expense of an increased number of iterations for convergence. The coupled system is shown in Fig. 8.

The computationally intensive fluid dynamics simulation step is performed on the Grid. The user is requested to provide quality of service requirements (see Sec. 3.5.2), such as maximum solution time and cost, using the interface shown in Fig. 9.

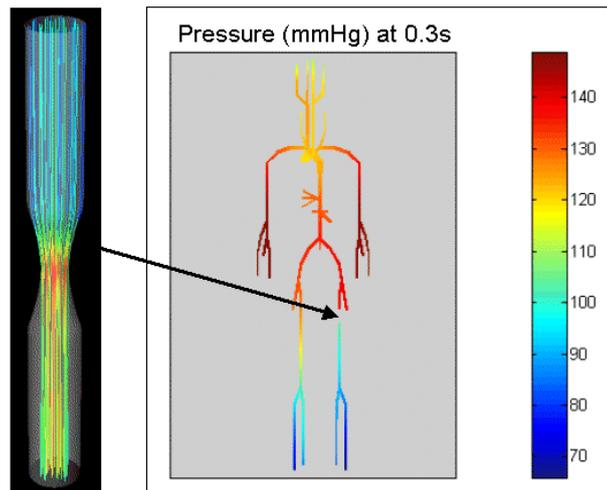


Fig. 10. Streamlines for 3D model and pressure distribution in compartments for peak flow in the stenosis

The processing time is dominated by solution of the 3D flow field, and the 3D model has 20967 nodes and 19264 hexahedral elements. For this model, approximately 80 minutes of grid time (8 processors) is required for one second of simulation time. When run locally, on a PC with a single processor (of similar performance to a Grid processor), solution takes three times longer. Results for the stenosis model are shown in Fig. 10.

2.3.3. Application to aortic arch geometry obtained from MR scan

Flow simulation of idealized geometries can provide valuable information for the design of vascular prostheses. However, in order for cardiovascular simulation to

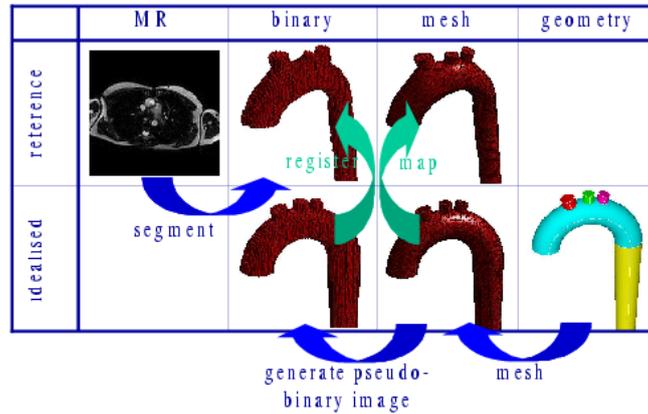


Fig. 11. Steps involved in generating the mesh

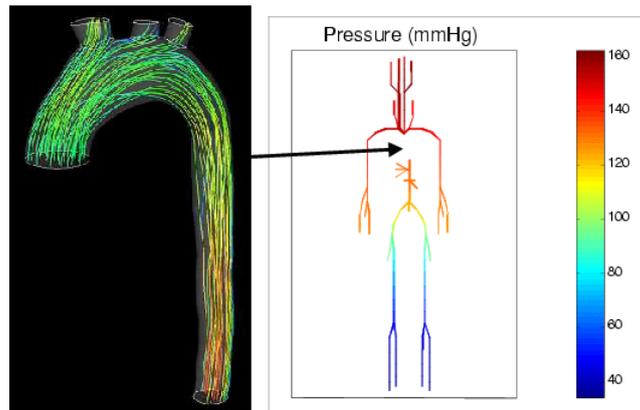


Fig. 12. Streamlines for 3D model and pressure distribution in compartments at peak systolic flow

be able to influence individual patient management, we need to be able to rapidly generate meshes of real arterial geometries. For the aortic arch we generated a generic geometric model from half a torus for the arch itself, a conical frustum for the descending aorta, and three cylinders for the major arteries leaving the arch. This geometry was meshed, and then a binary image created from the mesh. The MR scan was segmented manually to create a second binary image. The two binary images were registered, and the resulting mapping was applied to the geometric mesh to give a patient-specific mesh [Jones *et al.* 2004]. The process is illustrated in Fig. 11.

The patient-specific mesh was imported into the GEMSS cardiovascular soft-



Fig. 13. Wall displacement at systole (yellow) compared to diastole (red)

ware, and ‘Westerhof’ compartments coupled to the outlet boundaries. The solution at peak systolic flow is shown in Fig. 12.

The MR scan used in the above process was the first in a time series of seven scans. Using the image registration technique directly on subsequent MR scans in the series, it was possible to automatically obtain the vector mesh displacement field and hence vessel wall motion throughout the cardiac cycle, as shown in Fig. 13. The simulation results we obtained assume that the vessel wall is rigid, but it can be seen that the wall moves significantly during a cardiac cycle. In the future we intend to improve the fidelity of simulation by imposing the wall motion as a boundary condition on the 3D fluid dynamics simulation.

3. Part II: Grids for Medical Simulation Services

3.1. Introduction

In the first part of this paper, we have presented a broad selection of numerical simulation services for medical problems, which are all quite demanding in terms of computational resources, and in most cases go beyond what is available for specialized clinics, let alone a medical practitioner. A possible solution is Grid computing.

Grid computing aims to provide a global IT infrastructure for a coordinated, flexible, and secure sharing of diverse resources, including computers, applications, data, storage, networks, and scientific instruments across dynamic and geographically dispersed organizations and communities (cf. virtual organizations). Grids might be classified into Computational Grids, Data/Information/Knowledge Grids, and Collaborative Grids. Computational Grids, which are the focus of this paper, provide a platform for advanced high-performance and/or high-throughput applications that could not be tackled by a single system, or at least not by local systems, as in the case at hand. Data Grids, on the other hand, focus on the sharing of vast quantities of data. Information and Knowledge Grids extend the capabilities

of Data Grids by providing support for data categorization, information discovery, ontologies, and knowledge sharing and reuse. In the medical context, such Grids could be used to manage medical data sets such as images and the corresponding patient records. Collaborative Grids establish a virtual environment, which enables geographically dispersed individuals or groups of people to cooperate within virtual laboratories or to control and manage remote equipment, sensors, and instruments. A typical medical application could be tele-medicine or medical consulting by specialists.

However, the vision of large scale resource sharing has not yet become a reality in many areas. This can be attributed mainly to the lack of commonly accepted standards, as well as to the diversity and fragmentation of available Grid middleware, tools and services. According to a recent survey of twenty European Grid projects [Gagliardi *et al.* 2004], the most widely used middleware is the Globus [GLOBUS] toolkit followed by Unicore [UNICORE]. Recently, however, the Globus toolkit, which has been originally designed for the needs of HPC resource sharing in the academic community, has undergone a significant shift towards the adoption of a service-oriented paradigm, and the increasing support for and utilization of commercial Web Services technologies. The Open Grid Services Architecture (OGSA) [Foster *et al.* 2002] constitutes a significant effort in bringing Grid technologies and Web Services together. The fact that the implementation of OGSA will be based on the forthcoming Web Services Resource Framework (WSRF) [WSRF], currently standardized by OASIS [OASIS], is a further significant step in this direction and will allow the utilization of standard Web Services technologies, which enjoy large scale industry support, for Grid computing.

3.2. The GEMSS Grid

The GEMSS Project (Grid Enabled Medical Simulation Services) is mainly concerned with the computational aspect of Grids, focusing on developing a secure, service-oriented Grid infrastructure for the on-demand provision of advanced medical simulation services. GEMSS has demonstrated that Grid technologies can be used to enhance healthcare by bringing a variety of medical computing and resource services into the user's environment. The GEMSS Grid middleware can be used to provide medical practitioners and researchers with access to advanced simulation and image processing services for improved pre-operative planning and near real-time surgical support. The Grid also brings external computational resources to the medical technology industry, already using bio-numerics for virtual prototyping, but requiring larger-scale compute resources due to the growth in the complexity of design problems made tractable by numerical methodology advances.

Anticipating the shift of Grid technologies towards Web Services, the GEMSS Grid infrastructure and middleware adopts a service-oriented architecture built on top of standard Web Services technologies, ensuring future extensibility and interoperability. Key aspects of the GEMSS Grid middleware include negotiable quality-of-

service (QoS) support for time-critical service provision, flexible support for business models, and security at all levels in order to ensure privacy of patient data as well as compliance with the latest laws and EU regulations related to providing medical services over the Internet.

3.3. *Legal and Security Issues for Medical Simulation Services*

Most medical applications process medical image data that is impossible to anonymize. Because of this we are forced^a [Middleton *et al.* 2005] to treat all application data as personal data, and legally required to provide “best practice” security within a reasonable cost of implementation. The *controller* of the patient data, normally the hospital that generated it, must also have a contract with all *processors* and sub-processors (i.e. service providers), on the Grid before any data can be processed. The requirement for “best practice” drives our emphasis on security within the GEMSS project. A basic HTTPS security transport layer is provided for web service invocation, and built on top of this is support for the WS Security standard and a proprietary end-to-end security protocol. A project wide public key infrastructure (PKI) provides X.509 certificates that identify our users. This PKI underpins the WS Security, end-to-end security and authentication aspects within GEMSS. As a last line of security we have logging and an experimental intrusion detection system.

3.4. *The GEMSS Architecture*

The GEMSS architecture uses a client/server topology employing a service-oriented architecture. Since GEMSS supports a public key infrastructure there can be a number of trusted certificate authorities to issue security certificates. Service registries hold details of where services can be found and are queried before a client interacts with specific service providers. Fig. 14 shows this basic architecture.

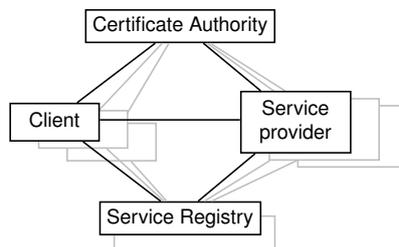


Fig. 14. High-level client/server architecture

The design that follows encompasses a three step process to job execution. First

^aat least in the EU, but similar rules should apply elsewhere

there is the initial business step, where accounts are opened and payment details fixed. The pricing model may also be chosen at this stage. Next there is a quality of service negotiation step, where a job's quality of service and price, if not subject to a fixed price model, is negotiated and agreed. Finally, once a contract is in place, the job itself can be submitted and executed. Fig. 15 shows this three-step process.

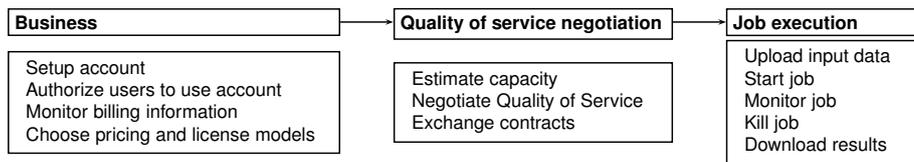


Fig. 15. Three step process to using the Grid

The client architecture is shown on the left hand side of Fig. 16. The client would typically run an offline business workflow to open an account with a service provider and agree a payment mechanism. When a Grid job is run the client would then open negotiations with a set of service providers for a particular application's job. A quality of service negotiation would then be run to request bids from all interested service providers who can run the clients job; this would result in a contract being agreed with a single service provider and a scheduler reservation made with an appropriate quality of service level. The client would then upload the job input data to the service provider where the server side application scripts would take over. The client always drives the three-steps of the GEMSS process. The reason for this is the requirement to operate with firewalls and not tunnel holes through them. More details on the client side infrastructure can be found in [Benkner *et al.* 2005].

The service provider architecture is shown on the right hand side of Fig. 16. Applications are run by the service under the direction and orchestration of the client, subject to the necessary availability of and authority to use resources. The compute resources available at the service provider's site will be used to actually run the services. The quality of service management module is there to handle reservations with the scheduler and provide input to the quality of service negotiation process so that sensible bids can be made to client job requests. The process of providing numerical simulation applications as Grid Services is described in more detail below.

3.5. *Grid-Enabling Medical Simulation Applications*

A major objective of the GEMSS project was the development of a generic Grid service provision framework that simplifies the task of transforming existing applications into Grid services without having to deal with the details of Web Services and Grid technologies.

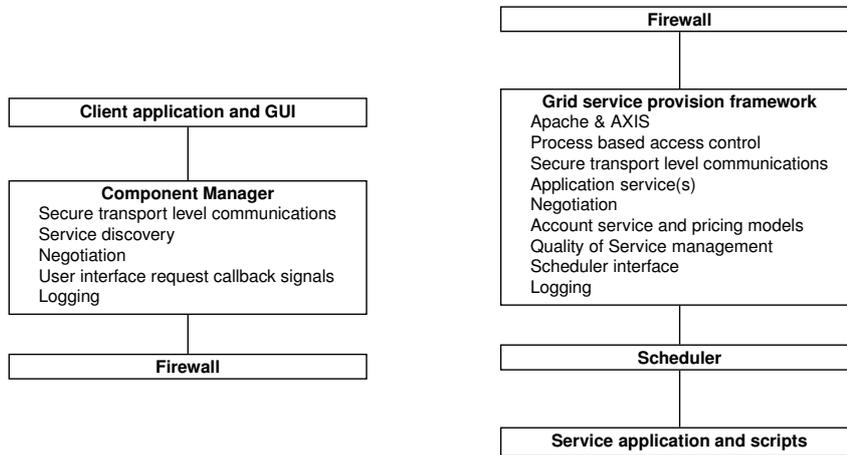


Fig. 16. Client architecture(left). Service provider architecture (right).

In order to Grid-enable a simulation application under the GEMSS environment it must be possible to partition the application between client and server. The client part is responsible for supplying input data, for controlling job execution and for processing of output data. The application kernel itself runs remotely as a batch job on some Grid host. Provided that an application can already be executed in batch mode, usually no code changes are required. This implies that in I/O operations files must not be accessed with absolute path names.

The GEMSS service provision environment requires that an application is pre-installed on some Grid host and a job-script to start the application is available. Moreover, if QoS support is required, a performance model to estimate the response time and a pricing model for determining the price of a job execution has to be provided (see below for more details).

3.5.1. Service Generation and Deployment

The transformation of medical simulation applications into Grid services is based on the concept of generic application services. A generic application service is a configurable software component which exposes a native application as a service to be accessed by multiple remote clients over the Internet. A generic application service provides common methods for data staging, remote job management, error recovery, and QoS support, which are to be supported by all GEMSS services. These methods are customized for a specific GEMSS application by means of an XML application descriptor which specifies the input/output files, the script for initiating job execution, and a set of performance-relevant application parameters required for QoS support. A generic application service is realized as a Java class,

which is transformed automatically into a Web Service with corresponding WSDL^b descriptions.

For exposing a simulation application as a GEMSS service, a corresponding deployment tool has been developed. The deployment tool enables the user to enter the information required in an application descriptor via a GUI and to control the deployment process. Internally the deployment tool creates the XML application descriptor, generates an appropriately customized Web Service which encapsulates the application, publishes the corresponding WSDL document in a registry service, and finally deploys the service within the GEMSS hosting environment. As a result, a simulation application is wrapped as a service and accessible over the Internet via SOAP^c.

3.5.2. *Quality of Service Support*

Quality of Service is a crucial issue in the context of providing medical applications on the Grid. The GEMSS Grid infrastructure offers application-level QoS support based on response time guarantees for the execution of time-consuming simulation services executing on remote Grid host. Response time guarantees are usually negotiated dynamically between a client and potential service providers on a case-by-case basis. Since GEMSS also addresses the realization of Grid business models (see [Benkner *et al.* 2005] for examples), services may be configured to support price negotiation as well. All explicitly formulated QoS guarantees such as execution time and price agreed between a service consumer and a service provider are expressed in form of an XML document following the Web Service Level Agreement (WSLA) [WSLA] specification. Besides explicitly negotiable QoS parameters, the GEMSS infrastructure provides implicit QoS by realizing highest security levels and providing support for error recovery.

Fig. 17 presents the basic parts of the GEMSS QoS Infrastructure separated into client-side and service-side parts. On the client-side a QoS component is provided which is utilized by client applications for QoS negotiation. The QoS infrastructure on the service side is centered around the QoS management module, which receives a QoS request from a client, checks whether the client's QoS constraints can be met, and returns a corresponding QoS offer.

The *application performance model* (APM) is utilized to estimate the performance requirements (runtime, memory and disc requirements) for a service request. The performance model takes as input a request descriptor, containing request meta-data (e.g. mesh size), and a machine descriptor containing information about available computational resources (e.g. number of processors, etc.), and returns a performance descriptor describing the estimated performance requirements. The *compute resource manager* (CRM) module provides a high level interface to an un-

^bWeb Service Definition Language, see www.w3.org/TR/wsdl

^cSimple Object Access Protocol, see www.w3.org/TR/soap

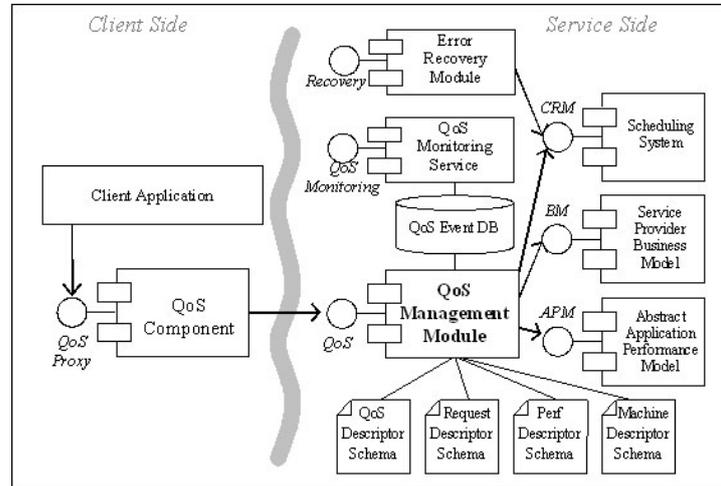


Fig. 17. GEMSS QoS Infrastructure

derlying scheduling system which has to provide support for advance reservation (e.g. NEC's COSY scheduler [Zimmermann and Cao 2004]). The *business model* (BM) module defines a generic mechanism to calculate a price based on estimated resource allocation. The *QoS monitoring service* is an independent service for monitoring all events of the QoS management module, and storing them in the QoS event database. The error recovery module supports check pointing and restart, if supported by the native application and the compute resource manager.

The basic QoS negotiation in GEMSS is based on a request-offer model where the client requests offers from service providers. If the client agrees to an offer, it is confirmed by the client and signed by both parties resulting in a QoS contract. A basic QoS negotiation process usually comprises the following steps:

- (1) In an initial task the client may access a GEMSS registry service to obtain a list of potential service providers to be contacted during the basic QoS negotiation.
- (2) Then the client generates a request descriptor and a QoS request which are sent to the potential service providers. The request descriptor contains meta data about the client request (e.g. mesh size, resolution, etc.) and the QoS request specifies the required QoS constraints (maximal cost, earliest start of job execution, latest finish time).
- (3) On the server side the QoS Management module attempts to generate a QoS offer by executing the performance model and business model and by checking with the compute resource manager if the required resources can be made available.
- (4) Based on the performance estimate and available resources, the QoS Management module makes a temporary resource reservation with a short expiration time and returns a QoS offer to the client.

- (5) On the client side, the QoS offers from different service providers are received and analyzed.
- (6) The client then confirms the best offer, or, if it is not satisfied with the offered QoS constraints, may set up a new QoS request with different constraints and continue with step (1).
- (7) On the server-side the QoS management module confirms the temporary resource reservation made for the offer, signs the QoS contract and returns it to the client.

After the basic QoS negotiation the regular job-handling workflow is initiated. This usually comprises uploading of input data, starting of the native application, and downloading of results. Within the GEMSS project also more sophisticated negotiation strategies based on auction models have been realized, a description of which is beyond the scope of this paper, see [Benkner *et al.* 2005] for details.

3.5.3. *Practical Experience with the GEMSS Grid*

During the GEMSS project, extensive experience with computational Grid services could be collected. First of all, such a Grid system (together with appropriate client software) has proved to effectively enable a user who is unexperienced with respect to high-performance computing to use such systems and applications. It also acts as a layer of abstraction on top of individual HPC environments, schedulers and policies, which often differ considerably in practice. The availability of several service providers offering the same service can greatly enhance the availability of the service, especially in connection with a service discovery component.

In general, only an application-specific performance prediction permits sensible scheduling policies of such a Grid system. The difficulty of such a performance estimation varies between applications. For instance, in the case of the SPECT image reconstruction service, a very accurate prediction is possible. On the other end of the scale, for non-linear CFD simulations and unknown geometries, it is quite difficult to estimate performance *a priori*. Here we have to use substantial safety margins to ensure a low job failure rate due to insufficient computing time. A possible future extension could support checkpointing and restarting an application. The “best” solution also depends on the business models used, for instance whether time not used (but reserved) has to be paid, or if simply a price per job is fixed.

The coordination of the different timing constraints during QoS negotiation and scheduling has turned out to be an interesting challenge. If time slots are too short, many negotiations will fail, if they are too long, negotiation will take too much time. While a number of problems have been fixed already, there still is room for optimization.

A basic integration of an application into the GEMSS system is not very difficult. On the client side, the necessary input has to be packed, and in principle a standard script can be used to start a job submission, including service discovery

and negotiation. Of course, more sophisticated user interaction is possible, cf. Fig. 9, p. 14. On the server side, only very modest requirements are posed on the behavior of the software, such as the ability to work with relative paths. The only essential new piece is the development of a performance estimation component. Here one often can start with a very simple and conservative estimation, which can be refined subsequently.

4. Conclusions

With a subjective selection and discussion of medical simulation applications we have tried to underline our conviction that such simulations have a huge potential for improving medical practice. Evidently, their state of maturity is very different. While some, like cardio-vascular system simulation or maxillo-facial surgery planning, are still in a research stage, others, like improved SPECT image reconstruction or dose-distribution calculation by Monte-Carlo methods, are in principle ready for the clinical routine. In general, however, we feel that the “golden age” of medical simulation is yet to come. What are the steps to be taken before simulation becomes medical practice?

With the availability of Grid systems like GEMSS, the practical accessibility of the necessary high-performance environments to clinical users is solved in principle, including aspects like security and lawfulness (the latter with respect to EU legislation). This marks an essential step forward with respect to prior state. However, the corresponding technology is fresh and not yet standardized, the necessary infrastructure (like certificates and certification authorities) not in place or not widely used.

Many medical simulation applications still lack a full clinical evaluation. While such trials are certainly mandatory for drugs and medical devices rather than for mere planning tools, it will be necessary to improve confidence into these methods, and they will have to prove that they lead to better results than current “simulation-less” practice.

Our work reinforced our conviction that the ergonomics of the user interface are crucial for the adoption of a medical application. This requirement looks quite different if the user is a clinician rather than an engineer or even a scientist: For instance, the term “boundary conditions” in general does not mean much to a physician, and so, the interfaces of commercial simulation packages are certainly not adequate. But ergonomics go far beyond the visible interface. In the section on maxillo-facial surgery planning (2.2), we discussed automation and robustness of the whole tool chain, which relies on guaranteeing certain topological properties of the input models. This can work only if expert knowledge is fed into a specialized simulation scenario; in our case, we use landmarks together with the knowledge that we operate on a head and that the mandible should be a separate bony component.

This specialization evidently requires considerable work; the open question is how much additional work will be required for adapting an application to new

26 REFERENCES

scenarios, a question that will ultimately decide the the commercial prospects of such applications.

When looking to the future of medical simulation research, we expect that parts that are separate today will become more and more integrated, led by the vision of a “virtual human” which is a complete biomechanical model of the human body. Using such a holistic approach, individual simulations could be provided with more realistic boundary conditions, and coupling of different phenomena would easier.

In the introduction, we contrasted bio-informatics (genomics and proteomics) and bio-numeric (organ-level simulation). In the future, we may expect multi-scale approaches which integrate different length- or time-scales, eventually blurring the distinction between bio-informatics and bio-numeric. However, the vision of a simulation “from molecules to man” will remain a vision for quite some time.

Acknowledgments

The projects SimBio and GEMSS project were funded by the European Commission under Grants IST-1999-10378 and IST-2001-37153. It is a pleasure for us to acknowledge the great collaboration and support of our colleagues in these projects. Images were kindly provided by A. McCarthy, University of Sheffield (Fig. 4) and W. Backfrieder & M. Forster, University of Vienna (Fig. 3). Data for Fig. 7 was provided by T. Hierl, University clinics Leipzig.

References

- Backfrieder W., Forster M., Engelbrecht G., and Benkner S. (2003). Optimized design of VOR for 3D image reconstruction in SPECT in a service oriented parallel implementation. *Journal of Nuclear Medicine Technology*, 31(2):139.
- Barber D. C. and Hose D. R. (2005). Automatic segmentation of medical images using image registration: diagnostic and simulation applications. *Journal of Medical Engineering & Technology*. In press.
- Benkner S., Berti G., Engelbrecht G., Fingberg J., Kohring G., Middleton S. E., and Schmidt R. (2005). GEMSS: grid infrastructure for medical service provision. *Methods of Information in Medicine*, 44(2). In press.
- Berti G. (2004). Image-based unstructured 3d mesh generation for medical applications. In Neittaanmäki P., Rossi T., Majava K., and Pironneau O., editors, *European Congress on Computational Methods in Applied Sciences and Engineering ECCOMAS 2004*. University of Jyväskylä, Department of Mathematical Information Technology. CDROM.
- Chabanas M., Payan Y., Marécaux C., Swider P., and Boutaul F. (2004). Comparison of linear and non-linear soft tissue models with post-operative CT scan in maxillofacial surgery. In *Medical Simulation: International Symposium, ISMS 2004*, number 3078 in LNCS, pages 19–27. Springer.
- de Hart J., Peters G. W. M., Schreurs P. J. G., and Baaijens F. P. T. (2003).

- A three-dimensional computational analysis of fluid structure interaction in the aortic valve. *J Biomechanics*, 36:103–112.
- Fenner J. W., Mehrem R. A., Ganesan V., Riley S., Middleton S. E., Potter K., and Walton L. (2005). Radiosurgery planning supported by the GEMSS Grid. *Studies in Health Technologies and Informatics*. In press.
- Formaggia L., Gerbeau J. F., Nobile F., and Quarteroni A. (2001). On the coupling of 3d and 1d navier-stokes equations for flow problems in compliant vessels. *Comput. Methods Appl. Mech. Engrg.*, 191:561–582.
- Foster I., Kesselman C., Nick J., and Tuecke S. (2002). The physiology of the grid: An open grid services architecture for distributed systems integration. *Open Grid Service Infrastructure WG, Global Grid Forum*.
- Gagliardi F. et al. (2004). IST grid projects inventory and roadmap. www.gridstart.org.
- GEMSS (2002). The GEMSS project: Grid-enabled medical simulation services. <http://www.gemss.de>. EU IST project IST-2001-37153, September 2002 – February 2005.
- Gladilin E. (2003). *Biomechanical Modeling of Soft Tissue and Facial Expressions for Craniofacial Surgery Planning*. PhD thesis, Fachbereich Mathematik und Informatik, Freie Universität Berlin. URL <http://www.diss.fu-berlin.de/2003/143/indexe.html>.
- GLOBUS. The GLOBUS toolkit. www.globus.org.
- Hose D. R., Lawford P. V., Narracott A. J., et al. (2003). Fluid-solid interaction: benchmarking of an external coupling of ANSYS with CFX for cardiovascular applications. *Journal of Medical Engineering & Technology*, 27:23–31.
- Hose D. R., Narracott A. J., Penrose J. M. T., Baguley D., Jones I., and Lawford P. V. (2005). Fundamental mechanics of aortic heart valve closure. *Journal of Biomechanics*. In press.
- HyPre (2004). The HyPre solver library homepage. www.llnl.gov/CASC/hypre.
- Jones D. M., Hose D. R., Lawford P. V., Hill D. L. G., Razavi R. S., and Barber D. C. (2004). Creation of patient-specific CFD models by morphing a previously-meshed reference geometry using image registration. In *Proceedings of Medical Image Understanding and Analysis, London, 2004*, pages 173–176.
- Kass D. A. (2002). Age-related changes in ventricular-arterial coupling: path physiologic implications. *Heart Fail Rev*, 7(1):51–62.
- Koch R. (2001). *Methods for Physics Based Facial Surgery Prediction*. PhD thesis, Institute of Scientific Computing, ETH Zürich. Diss. No. 13912.
- Lagana K., Balossino R., Migliavacca F., Pennati G., Bove E. L., de Leval M. R., and Dubini G. (2005). Multiscale modeling of the cardiovascular system: application to the study of pulmonary and coronary perfusions in the univentricular circulation. *Journal of Biomechanics*. In press.
- Lippmann H. and Kruggel F. (2005). Quasi real-time neurosurgery support by mr image processing via grid computing. *Neurosurgery Clinics of North America*, 16

28 REFERENCES

- (1):65–75.
- McCarthy A. D., Wilkinson I. D., Hose D. R., Barber D. C., Wood S., Darwent G., Chan D., and Bickerstaff D. R. (2002). Musculo-skeletal simulation: Finite element meshes derived from magnetic resonance volumes. In *Proceedings of the International Society for Magnetic Resonance in Medicine, May 18-24, 2002*.
- Middleton S., Herveg J., Crazzolara F., Marvin D., and Pouillet Y. (2005). GEMSS: Security and privacy for a medical grid. *Methods of Information in Medicine*. To appear.
- Migliavacca F., Balossino R., Pennati G., et al. (2004). Multiscale modelling in biofluidynamics: application to reconstructive paediatric cardiac surgery. In *14th European Society of Biomechanics Conference, s’Hertogenbosch, The Netherlands, July 2004*.
- OASIS (1993). OASIS - organization for the advancement of structured information standards. www.oasis-open.org.
- Quarteroni A. and Venezian A. (2001). Analysis of a geometrical multiscale model based on the coupling of ODEs and PDE for blood flow simulation. *Multiscale Model. Simul.*, 1(2):173–195.
- Schmidt J. G., Berti G., Fingberg J., Wollny G., and Cao J. (2004). A finite-element based tool chain for the planning and simulation of maxillo-facial surgery. In Neittaanmäki P., Rossi T., Majava K., and Pironneau O., editors, *European Congress on Computational Methods in Applied Sciences and Engineering EC-COMAS 2004*. University of Jyväskylä, Department of Mathematical Information Technology. CDROM.
- Schutyser F., Cleynenbreugel J. V., Schoenaers J., Marchal G., and Suetens P. (1999). A simulation environment for maxillofacial surgery including soft tissue implications. In *Proceedings of MICCAI 1999*, pages 1210–1217.
- Segers P., Stergiopoulos N., Verdonck P., and Verhoeven R. (1997). Assessment of distributed arterial network models. *Journal of Medical and Biological Engineering and Computing*, 35:729–736.
- Segers P. and Verdonck P. (2002). *Theory and practice of vascular medicine*, chapter Integrated Physiology of the Vascular System, pages 116–137. Springer Verlag. ISBN 3-540-41484-3.
- Simbio (2000). The SimBio Project. <http://www.simbio.de>. EU IST project No. 1999-10378, March 2000 – June 2003.
- Steinman D. A. (2002). Image-based computational fluid dynamics modeling in realistic arterial geometries. *Annals of Biomedical Engineering*, 30:483–497.
- Stergiopoulos N., Westerhof B. E., and Westerhof N. (1998). Physical basis of pressure transfer from periphery to aorta: a model-based study. *Am J Physiol*, 274(4 Pt 2):H1386–92.
- Stergiopoulos N., Westerhof B. E., and Westerhof N. (1999). Total arterial inertance as the fourth element of the windkessel model. *Am J Physiol*, 276(1 Pt 2):H81–8.
- UNICORE. UNiform Interface to COmputing REsources. www.unicore.org.

- Westerhof N., Bosman F., Vries C. J. D., et al. (1969). Analog studies of the human systemic arterial tree. *Journal of Biomechanics*, 2:121–143.
- Wolters C. H., Kuhn M., Anwander A., and Reitzinger S. (2002). A parallel algebraic multigrid solver for finite element method based source localization in the human brain. *Computing and Visualization in Science*, 5(3):165–177.
- WSLA (2003). Web service level agreement (WSLA) language specification. www.research.ibm.com/wsla/WSLASpecV1-20030128.pdf.
- WSRF (2005). Web services resource framework (WSRF). www.oasis-open.org/committees/wsrff.
- Zachow S., Gladilin E., Zeilhofer H.-F., and Sader R. (2001). Improved 3D osteotomy planning in cranio-maxillofacial surgery. *Lecture Notes in Computer Science*, 2208:473–481.
- Zhao S., Xu X., et al. (1998). The numerical analysis of fluid-solid interactions for blood flow in arterial structions. Part 2: Development of coupled fluid-solid algorithms. *Proceedings of the Institution of Mechanical Engineers Part H- Journal of Engineering in Medicine*, 212(H4):241–252.
- Zimmermann F. and Cao J. (2004). Queue scheduling and advance reservations with COSY. *International Parallel and Distributed Processing Symposium, Santa Fe, New Mexico*.