The IST Programme
Project No. IST-2001-37153

GEMSS

Grid-enabled Medical Simulation Services

http://www.gemss.de

Deliverable D6.2a
First Annual Project Progress Report

Status:	Release
Version:	1.00
Security:	**Public**
Responsible:	NEC – Jochen Fingberg
Authoring Partners:	All Partners

Release History

Version	Date	
0.1	04.08.2003	Initial Draft
0.2	11.08.2003	2nd Draft
0.3	13.08.2003	3rd Draft
0.4	14.08.2003	4th Draft
0.5	14.08.2003	5th Draft
0.6	15.08.2003	6th Draft
0.7	18.08.2003	7th Draft
0.8	18.08.2003	8th Draft
0.9	18.08.2003	9th Draft
0.10	19.08.2003	10th Draft
0.11	20.08.2003	11th Draft
0.12	21.08.2003	12th Draft
0.13	22.08.2003	13th Draft
0.14	26.08.2003	14th Draft
0.15	27.08.2003	15th Draft
0.16	27.08.2003	16th Draft
0.16	27.08.2003	16th Draft
1.00	29.08.2003	1st Release

The GEMSS Consortium:
NEC Europe Ltd. – UK
ISS – Austria
CRID – Belgium
ASD – Germany
IT-Innovation – UK
MPI of Cognitive Neuroscience – Germany
IBMTP – Austria
CFX (AEA) – UK
IDAC – Ireland
Sheffield University – UK

First Annual Project Progress Report

Executive Summary:

This annual report covers project months 1-12 (September 1st, 2002 - August 31st, 2003) of the GEMSS project and provides a summary of technical progress made. Activities reported are in line with the project workplan as covered by the GEMSS Contract IST-2001-37153. The First Annual Project Progress Report will be available as a public document.

Table of Content:

1 Project Overview

GEMSS is developing an interoperable, innovative Grid middleware for medical service applications building on common Grid standards. The focus is on innovative extensions that support medical applications including security models compliant with European legal issues, fail-over and recovery from errors as well as workflow and service orchestration techniques for time-critical processes. GEMSS is a two and a half year project which started in September 2002.

1.1 Objectives

The central objectives of the GEMSS project are to:

- demonstrate that the Grid can improve pre-operative planning and near real-time surgical support by providing access to advanced simulation and image-processing services,
- build middleware on existing/developing Grid technology standards to provide support for authorisation, workflow, security and Quality of Service aspects,
- develop, evaluate and validate a test-bed for the GEMSS system, including its deployment in the end-user's working environment,
- analyse and test the European Legal Framework from three approaches (Privacy, Contract, Liability) considering the use of Grid-based-technologies / applications partially provided through the Internet, in order to know if the European Legal Framework permits the development and the exploitation of such applications and if positive, under which constraints.

1.2 Main goals

The main goal of GEMSS is to provide end-users from the medical community with advanced tools at their workplace through easy-to-use interfaces. In particular GEMSS will:

- install a secure, extendible, interoperable and collaborative test bed for GRID-enabled medical application services,
- demonstrate the medical significance of the GEMSS models,
- demonstrate the functionality of the GRID-infrastructure,
- open new business models for future commercial exploitation,
- highlight and show the main legal aspects to be considered when developing and implementing GRID-based technologies.

1.3 Priorities

Through the central objective outlined in section 1.1 GEMSS is an infrastructure project driven by its 6 prototype medical service applications.

1.4 Architecture

The GEMSS architecture is based on a service oriented client-server topology. Figure 1 shows the initial basic design.

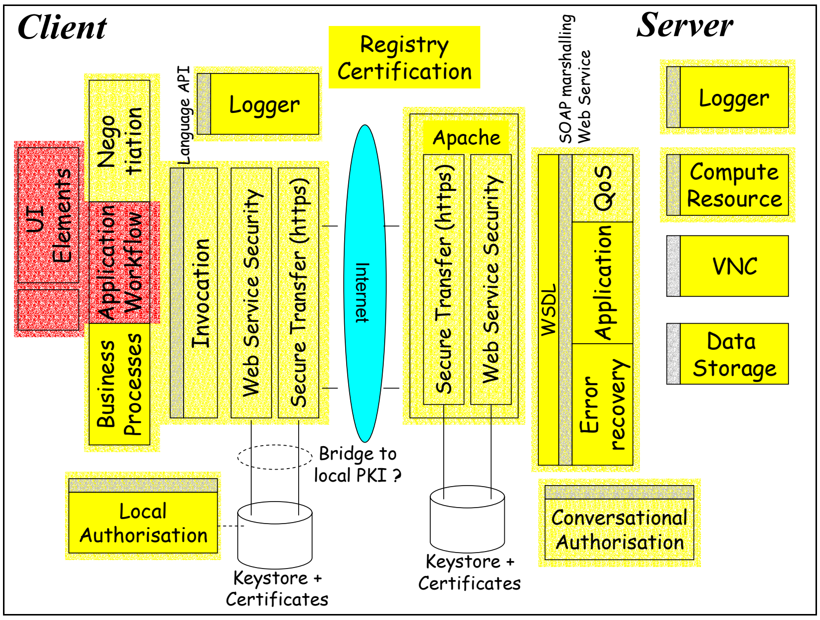

Figure 1: Initial Design of the Global GEMSS Architecture

2 Consortium

The Consortium consists of 10 partners, with three experienced in the GRID, two software developers, two University medical departments, and three end users, two of whom are specialist SME's in the area of bio-medical simulation. The Consortium thus represents a well balanced mixture of private and public partners whose activities range from basic research to industrial/commercial development and sales. The heterogeneous character of the Consortium has been designed to form a critical mass to best address the scientific requirements imposed by the problems to be solved.

No	Partner	Country	Specific Expertise, Role in the project
1	NEC	UK	Finite element simulation, mesh generation, GRID-middleware, HPC software & hardware, Project Management.
2	MPI	Germany	Medical image processing, Model validation, Acquisition of scan-data, medical end-users.
3	ISS	Austria	Programming environments, software tools, HPC tools and hardware, GRID technology.
4	USFD	UK	FSI modelling, cardiovascular and respiratory applications. Provision of validation datasets, scan data. Clinical test site for radiosurgery application. Medical end user.
5	AEA (1)	UK	Software and applications.
6	IT-Innovation	UK	GRID technology and know-how, technology transfer.
7	CRID	Belgium	Consultant in legal issues

8	ASD	Germany	FE and CFD consultancy for artificial organs and biomedical devices.
9	IDAC	Ireland	Stress-analysis consultancy, Consultancy, ANSYS re-sellers, GUIs
10	IBMTP	Austria	Image reconstruction software, acquisition of scan data, clinical end users
11	ANSYS-CFX (2)	UK	Software and applications.

(1) until 28.2.2003.
(2) from 1.3.2003.

The following table shows the involvement of partners in work packages. Work package and sub-task leaders are marked in red.

Partner	ST1.1	ST1.2	ST1.3	ST2.1	ST2.2	ST3.1	ST3.2	ST3.3	ST4.1	ST4.2	ST4.3	ST4.4	ST4.5	ST4.6	WP5	WP6
NEC	●	●	●	●	●	● (red)	● (red)	●	● (red)						●	● (red)
MPI	●		● (red)						●	● (red)					●	
ISS	●	●	●	● (red)	●	●	●	●						●	●	
USFD	● (red)		●								● (red)	● (red)	● (red)		●	
CFX/AEA	●		●	●		●	●					●	●		● (red)	
IT-Innov.	● (red)	● (red)	●	● (red)	● (red)		●				●				●	
CRID	●		●		● (red)										●	
ASD	●		●									●	●		●	
IDAC	●		●	●		●	●	● (red)					●		●	
IBMTP	●		●											● (red)	●	

3 Progress Report

3.1 Workpackage 1: System Design and Evaluation

This workpackage has three main responsibilities:

- to capture the system and end-user requirements,
- to produce a global design for GEMSS including the medical service applications and
- to evaluate the testbeds and the medical service applications within the testbeds produced during the course of this project.

3.1.1 Sub-task 1.1: Requirements Capture

A requirements capture process has been undertaken, which involved all GEMSS partners and culminated in the release of deliverable D1.1 “GEMSS requirements”. Specific partner requirements were first elicited, from both the technical and medical partners, and then unified to produce a set of project infrastructure requirements and application specific requirements. The project infrastructure requirements drive the Grid infrastructure design and implementation. The application specific requirements drive the medical application’s design and implementation. In addition to project and application requirements, existing end user resources were identified and existing application use cases drafted. Deliverable D1.1 provided a sound basis for the system design and WP4 application development.

3.1.2 Sub-task 1.2: System Design

A design process was undertaken involving primarily the technical partners and to a lesser extent the medical end-users. An incremental design methodology was chosen for the project to allow lessons learnt throughout the project lifetime to be fed back into the design and implementation through a series of development phases. A service-oriented client / server architecture was chosen, and a set of Grid use cases created. These were used to identify the infrastructure operations required and thus allowed specific functionality to be identified and system module interfaces defined. An application API was specified to provide guidance to the medical end users as they develop their Grid-enabled applications, ready for future integration with the GEMSS Grid infrastructure. Deliverable D1.2a "Global system design" was released as a result of this sub task. Deliverable D1.2a is a living design framework for GEMSS implementation work, which will be incrementally updated as the project progresses through the set of development phases.

3.1.3 Sub-task 1.3: Evaluation

The purpose of ST1.3 is to define the development methodology, to evaluate the series of testbeds developed in the projects and refine the requirements. Metrics for the robustness, scalability and performance will be defined to ensure that evaluation covers all aspects of the system. The evaluation will provide input into the refinement of requirements and design.
For the midpoint evaluation at PM 18 are planned:

- a survey of the user and developer experience with the GEMSS architecture,
- a re-evaluation of use case diagrams and system goals in light of user and developer comments, and
- an evaluation of software development progress in terms of system design metrics.

Since during PY 1, the GEMSS architecture was not available in a prototype stage, take-up of actions for ST 1.3 is planned for September, 2003, at PM 13.

3.2 Workpackage 2: Grid Services & Security

This workpackage is responsible for producing tools aimed at providing quality of service guaranties and for providing input on security issues, including answers to questions about legal issues for Grid-based medical services with regard to European legal provisions.

3.2.1 Sub-task 2.1: Workflow and Quality of Service

Workflow:
Initial detailed design work, followed by implementation, has been undertaken on the basic Grid infrastructure. This is required before workflow enactment can occur.

Detailed designs for two components, the "security context" and "transport and security" components, have been drafted. These components provide the basic job handling functionality. Implementation is underway for these two modules, in addition to the required modifications to the Apache web server required to support secure web service invocation.

The transport and security module provides other GEMSS modules with support for invoking WSDL based Grid services within a security context supplied by the security context module. The WSDL based Grid invocation includes the following:

- Support for SOAP-RPC encoding of WSDL messages
- Support for SOAP-Document encoding of WSDL messages
- Support for applying different security encoding to WSDL messages
- Support for applying different security policies for a given security encoding to WSDL messages.

The implementation target, for the first implementation phase I1, is to support the following:

- Invocation of standard WSDL web services across non-secure HTTP, where operation arguments are not complex.
- Invocation of standard WSDL web services across secure HTTPS, where operation arguments are not complex.
- Invocation of standard WSDL web services across HTTP and HTTPS where operation arguments are complex and a member of the consortium has produced a suitable Axis Serialiser and Deserialiser for each complex type.
- Secure invocation of WSDL web services using WS-Security encoding and a fixed policy concerning message security. The fixed policy is likely to only state that the message body should be signed. This security can be used across HTTP or HTTPS.

On the server side there will be a Web Hosting Platform based on Apache HTTP server that will probably support the following:

- Verification of WS-Security encoded SOAP messages using the Security Context Module available to the Web Hosting Platform
- Creation of WS-Security encoded SOAP response messages using a static policy whereby the SOAP message body is signed in combination with the Security Context Module.

Quality of Service:
This subtask aims at the design and development of an infrastructure capable of the provision of Quality of Service support for GEMSS application services. During the requirements capture phase, ISS led the definition of different QoS measures for each application. Based on the identified QoS requirements, a high level design of a QoS support infrastructure has been devised (see Deliverable D1.2a "Global system design").

The high level design of the QoS support infrastructure was presented at a project meeting held in March and subsequently agreed upon by the partners. A detailed design of a generic QoS API was specified by ISS and agreed with NEC and IT Innovation at the technical meeting in Dublin on June 10th-11th.

Moreover, ISS evaluated the performance of different web services technologies and frameworks (SOAP data transfer, SOAP security, etc.).

3.2.2 Sub-task 2.2: Security and Legal Issues

Security:
The security issues involved in the GEMSS project have been examined and a security methodology created and released as deliverable D2.2a "Methodology for Assessing

Security”. Potential threats are detailed, likely assets described, attack types listed and defence options identified, all in the context and scope of a GEMSS Grid. A security assessment is only meaningful, however, when performed on a specific partner site with full knowledge of the underlying local network infrastructure, security procedures and personnel. Deliverable D2.2a provides the method to perform such an analysis, but individual partners must actually perform the analysis for themselves.
Work has also been conducted within this sub task to disseminate security know-how within the project, and influence the requirements capture and design sub tasks. As such the global system design for GEMSS includes various security features as an integral part of its architecture, rather than added on after the fact. Work has also been performed to alert the medical partners of security issues within their application code, prior to integration with the future GEMSS Grid infrastructure.

Legal Issues:
The European Privacy Legal Framework has been analysed considering the development and the exploitation of the six specific applications of the GEMSS project. The main result is that these applications may be developed and exploited under certain conditions. It was demonstrated that the service provider is a processor (sub-contractor) of the data processing’s controller. The controller and the processor must be legally or contractually bound (it implies that only well-identified partners operate the applications). The medical data should be processed by a health professional subject to professional secrecy, *or* by another person subject to an equivalent obligation of secrecy, *or* by individuals or bodies working *on behalf* of a health care professional (like a processor) and also subject to such secrecy obligations. These results are relevant and may be extended to all similar GRID-based applications in the health care sector.

3.3 Workpackage 3: Medical Simulation System

This workpackage has the main task of providing access mechanisms for the GEMSS end users and integrating all the components together into a single system. Before launching into the construction of GEMSS, this workpackage is investigating the different means of providing ergonomic end-user access.

3.3.1 Sub-task 3.1: Portals and Access

During the design phase it was decided to forgo the construction of a system wide portal solution in favor of a library having a high level API which the applications would embed in their own GUIs. At the project meeting held in March, this decision was discussed with all the project partners, who subsequently agreed to support such an approach.

For ST3.1, this decision means that its main activities are being directed toward the creation of a simple, high level API for the applications to use as well as the creation of a component execution environment in which the components and the applications can interact. The first design of the API and component framework was created by NEC and presented to ISS and IT Innovation at a technical meeting held in Sankt Augustin on July 29th-30th, where it received a positive reception along with suggestions for changes. A second design draft, incorporating these suggestions, was then posted to the projects BSCW server on August 15th. The first prototype of the API and component execution environment will be delivered at the beginning of September.

A second activity being pursued by ST3.1 is an evaluation of various Grid alternatives in light of the requirements produced by ST1.1. For this purpose NEC has produced an evaluation template which NEC, ISS and IDAC will use to evaluate Unicore, OGSA and Web Services respectively. The results of this evaluation will feed into the deliverable D3.1a due in November.

For the purpose of evaluating OGSA, ISS has been tracking OGSA developments since technology preview 3, whereby they have concentrated on the Java implementation coupled with Tomcat as the hosting environment, but have also evaluated other OGSA compliant environments like the JBoss EJB server. In support of this activity ISS has also been investigating issues involved with the deployment of the medical image reconstruction service as an OGSA service. This included developing a prototype that exploits certain OGSA specific features like ServiceData and Notifications. Another issue of their research is to evaluate different invocation mechanisms together with OGSA services (e.g. WSIF) to enhance interoperability.

For the purpose of evaluating Unicore, NEC has deployed a Unicore based Grid internally and has implemented the basic services needed for Maxillo-facial surgery simulation.

IDAC have set up a web based system which delivers stent simulation applications over the internet. EASA has been used for both front and back end management. The system can be accessed at www.idacapps.com/easa. For security reasons, this address is not publicly visible. IDAC's firewall must be pinholed with the enduser's IP address to allow access. IDAC have two stent manufacturers who are currently evaluating the stent applications over the internet and providing feedback both on the delivery mechanism (quality of service, security) and on the stent applications.

3.3.2 Sub-task 3.2: System Integration and Testbed Deployment

During the first phase of this sub task, where the integral Grid related server and client packages are still being under development, implementation work has been focused more on associated areas like resource management. Even though there strong interactions between server based services like QoS and the resource management system, the NEC scheduler COSY has been modified to drive the first prototype test bed planned for the review (see below). The preparations for the review will them address the main topic of this sub-task: system integration. At the time of writing this text, IT Innovation is preparing a document that deals with the prototype requirements for both the application and the grid infrastructure site.

In the last technical planning meeting held on 29.-30.7.03 in Sankt Augustin, it was agreed to install a project-wide CVS repository at NEC to store the Grid related software packages. The overall version management for the integrated Grid software will be done by NEC.

Generic Application Service:
The Generic Application Service is a central component on the server side to expose native applications as web services. A first draft of the Generic Application Service Design was created by ISS and posted to the projects BSCW server on July 24th. A final Application Service specification for phase one of the implementation has been agreed with NEC and IT Innovation at the Implementation meeting on July 30th. A revised version of the Generic Application Service Design document with partner suggestions was posted to the BSCW server on August 28th server

COSY development for GEMSS:
The COSY scheduling system for PC clusters can be used as a resource manager in the GEMSS architecture. The common interfaces of the resource management module have been defined requiring key functionalities like advance reservation and starting jobs. The COSY development plan that will fit into the GEMSS implementation design is as follows:

1. Advance reservation

The user can specify explicitly the start time and runtime of the reservation. In this situation, required nodes will be allocated exactly at the required start time and last for the given runtime. The user can also specify explicitly the latest stop time and runtime of the reservation. In this situation, the reservation is firstly made to stop at the exact required time. After that the execution can be made earlier if there are available resources in a earlier time.

2. Two phase advance reservation

The user can do a general advance reservation as described in 1. The reservation will be effective for a certain period of time. Within the period, the user is required for an additional confirmation. Otherwise, the reservation will be released when expired.

3. Coupling with MPI environment

In order to execute the parallel program, COSY needs to integrate with existing MPI environment on the system. Current works include the integration of COSY with NEC and SCORE MPI implementations, which are available on the NEC cluster *grisu*.

4. Java interface definition

The current server side design of GEMSS foresees that both the actual application service and the QoS service interact with the resource scheduler. Since these entities are realized fully in

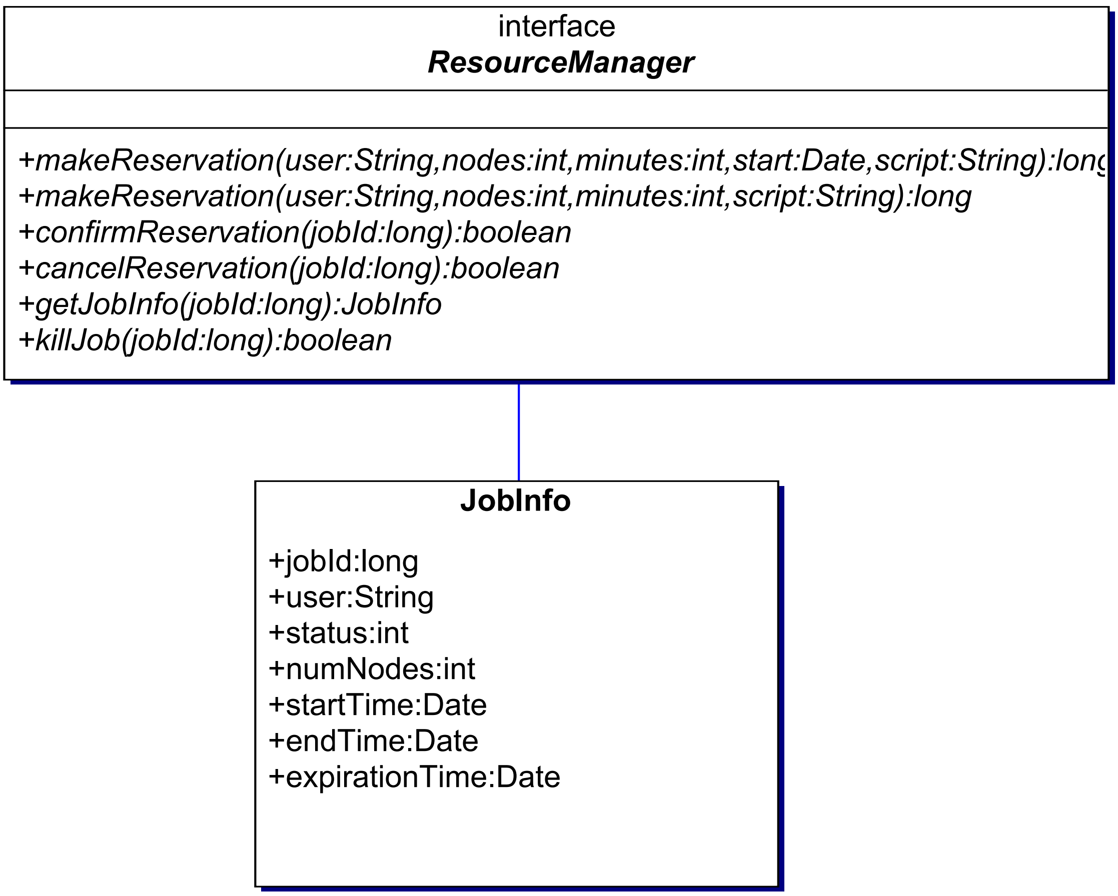

Java, we also define general Java interfaces for the scheduler. These interfaces are general in the sense that they are not perfectly tailored for one specific system (Cosy). We have to take into consideration that on different PC-clusters different scheduling systems are employed.

3.3.3 Sub-task 3.3: Grid-based Support and Consulting

The goals of this subtask are twofold:

1) To investigate Commercial Off The Shelf (COTS) collaborative tools and teleconferencing software.
2) Create recommendations as to how such software can be combined with consulting practices and integrated into the GEMSS system.

At the time of writing this subtask has been running two months. Point (2) can only be addressed when Point (1) has been fully investigated, so therefore IDAC has been mainly concerned with preparatory work for (1) above.

IDAC as a first step e-mailed all the partners using the GEMSS wp3@gemss.de mailing list for advice and sharing of experiences of COTS tools. IDAC also uploaded a draft document to the BSCW server which will form the backbone of the final report encompassing (i) and (ii) above. For each technology tested, IDAC will complete an analysis under each of the following headings:

- Installation
- Ease of use
- Cost
- Availability
- Compatibility
- Security
- GRID Features

IDAC can be seen in this context as an end-user of the GRID technology, so therefore IDAC has always a sharp focus on the commercial viability of any tests. In this vein IDAC finds that the most commercially-useful technologies are as follows:

- Desktop sharing applications, e.g. NetMeeting, VNC, GoToMyPC
- Application deployment across an intra or extranet, e.g. EASA.

To date, IDAC has added the following applications to its list:

Remote Desktop Sharing Applications:

- GoToMyPC
- WinVNC
- NetMeeting
- WebEx
- CuSeeMe
- VIC
- RAT
- CoCreate

Application Deployment Applications:

- EASA

Coupled with the fact that EASA and VNC both seem to be feasible to integrate with the GRID, investigating these two pieces of software further will be the main thrust of this subtask. IDAC also hope to investigate videoconferencing tools such as NetMeeting, CuSeeMe, VIC/RAT and other alternatives, however videoconferencing is of a lower priority than remote control and application deployment technologies, and also from discussing with some of the other partners, seem to be more difficult to tie into the GRID infrastructure.

IDAC are currently looking at setting up a VNC connection tunnelled over SSL on a Windows server as outlined in the document ***stunnel_vnc_for_windows*** on the BSCW server. We first looked at tunneling using SSH, however this is considered not secure enough for GRID application. Basically this involves using VNC connecting through stunnel (a program supporting the SSL protocol) between a client and a server, as in Figure 1 below.

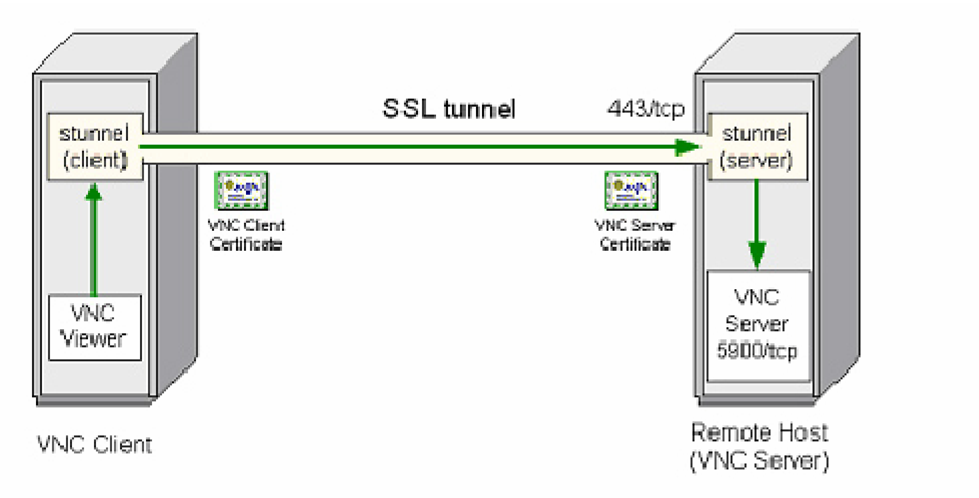

Figure 1: Typical SSL – VNC configuration

IDAC also currently have a couple of applications deployed on an EASA system and will as a part of this subtask, look at integrating one of these with the GRID, using EASA as a portal. This is dependent on EASA being made GRID-compatible in some way.

3.4 Workpackage 4: Medical Service Applications

This workpackage is responsible for the adaptation and integration of the medical simulation and image processing software into the GEMSS testbed. It includes a variety of medical service applications which

- have different GRID requirements concerning computation time (near real-time requirements vs. standard batch processing), memory usage, encryption, etc. and
- address different medical areas (cranial, pulmonary, cardio-vascular system) to address an end-users group of sufficient size for service provision.

3.4.1 Sub-task 4.1: Maxillo-facial Surgery Simulation

Background:
In patients, suffering from severe maxillary hypoplasia and retrognathia, conventional therapeutic surgery often fails to guarantee long-term stability. Using a rigid external distraction system for midfacial distraction osteogenesis is a new method to correct the underdevelopment of the midface, surpassing traditional orthognathic surgical approaches for these patients. Currently surgical planning is only based on CT images. The treatment consists of a midfacial osteotomy (*bone cutting*) followed by a halo-based distraction (*pulling*) step. The goal of this sub-task is the modeling of this distraction process to allow predictions on its outcome.

Tool chain Development:
The tools required to model and simulate this process perform such a simulation form a complex chain. First, the image data acquired (typically CT data) for a specific patient have to be converted into a format understood by all tools and potentially filtered to reduce noise effects like streak artifacts due to teeth fillings. Next, it has to be segmented into different tissue classes, which can be done either by simple intensity-based segmentation or later by registering to a detailed template model (*model-based segmentation*). A 16-bit CT reference image has been generated and segmented for this purpose together with a description of the relation between grey-value in the CT image and the material class in the resulting model. Based on the segmented augmented image the surgeon can specify cuts and distraction parameters using graphical tools. From the parameters output by these tools a geometric model of the surgical intervention is generated. The resulting process can be simulated by an FEM application. Finally, the results of the simulation have to be visualized and used to evaluate the quality of the treatment plan.

Workflow and Grid enabling:
This application can be characterized as distributed on-demand super-computing. The Triana workflow-based distributed problem solving environment that can be downloaded from: http://www.trianacode.org/projects/triana/ has been selected for the integration of the maxillo-facial tools. Once available, Triana will be interfaced to the GEMSS middleware to distribute the application workflow in a flexible way. Typical sizes of files to be transferred have been identified and the use of compression has been evaluated (factor of 50 for segmented images. Two basic scenarios for the Grid-use have been defined: an optional *thin client*, where the entire tool chain is executed on a server, and the client connects to the server using a remote control tool such as tightVNC (http://www.tightvnc.com); or a *thick client* where especially interactive visualization tasks run locally on the client side, except the final simulation run which needs High Performance Computing facilities. Concerning Quality of Service (QoS) requirements and error handling, the initial focus was on the FEM simulation, as it is both most challenging and time consuming. To predict its computational requirements and performance (memory, disk space, run-time, etc.) the application can provide a QoS vector consisting of typical problem parameters like mesh size, distraction length, model class or solver type to enable the prediction of resource requirements and performance (possibly based on previous runs). Different types of errors such as too large deformations decreasing mesh quality beneath a tolerable limit have to be anticipated. This can and will be detected from within the application. Another potential cause of failure is insufficient memory. If this is detected, there are several choices (automatic according to a pre-defined strategy or left to the user): Restart the application using a larger number of nodes (which will increase costs), choose a solver with lower memory consumption (which will increase the run-time), or reduce the model size (reduces accuracy). A third class of errors is failing convergence, which

may have occurred even if the program has terminated "normally" meaning that the problem is ill-conditioned. If just the necessary residual norm reduction has not been reached, the user can be offered to change relevant algorithmic parameters (higher maximal iteration counts). In the case of divergence, the user should contact the provider of the simulation program. The same applies to any other error (abnormal termination).

Sample Problems and Test Cases:
First an engineering example consisting of a square with a hard core (*bone*) and a soft surrounding (*soft tissues*) has been set-up as a benchmark and verification problem. Next, a medical test case consisting of a CT head scan image with added osteotomy lines specified by a surgeon, has been set-up to test all parts of the system.

Status Summary:
An initial version of the entire tool chain is running with some components requiring further enhancements after the next evaluation step. A basic environment has been set-up based on the following tools:

- Distributed workflow/GUI tool: Triana,
- An initial high-resolution CT-based head template model has been created,
- Basic registration of individual patient head is working,
- Model generation for one patient data set is working,
- First version of interactive cutting tool with multiple cut-lines is running under OpenDX (http://www.opendx.org),
- Boundary-adapted tetrahedral mesh generation operational, but so far limited to one material class (i.e. soft tissue, treating bone as background)
- Support for hyper-elastic and visco-elastic materials, geometric nonlinear FEM, prescribed forces and displacements,
- Basic tool chain integrated as a toolbox into the Triana workflow environment.

Next Steps:

- Integrate with GEMSS middleware,
- Optional: thin client with remote visualization using a secured version of tightVNC,
- Optional: Set-up and evaluate data base with measurements (run time, memory, disk space consumption, etc.) from previous runs for performance prediction purposes,
- Improve error handling and robustness of simulation software:
  Time step control, mesh quality monitoring, convergence monitoring, re-meshing if necessary,
- If necessary: Enhance template model with muscle fibers directions, and support anisotropic material models for FEM,
- Enhance cutting to allow visualization of connectivity,
- Derived E-module of bone from CT data,
- Possible extensions: Advanced mesh generation - support local refinement and multi-material marching tets.

3.4.2 Sub-task 4.2: Neuro-surgery Support

Background:
The major shortcoming of image-guided surgical planning based on pre-surgically acquired functional MRI (fMRI) data is the brain shift phenomenon. The occurrence of surgically induced deformations invalidates positional information about functionally relevant areas.

This problem is addressed by non-linear registration of pre-operative fMR images to intra-operative MRI acquired by an Open-MR scanner, or to intra-operative 3D ultrasound data.

Workflow and Grid Enabling:
As requirements for the chain to work, a pre-operative MR scan together with its aligned fMR data of the patient is needed. During the first stage of the surgery, before the skull is opened, an usually low-resolution image with the Open-MR scanner is acquired. After the correction of possible RF-Field inhomogeneities a linear registration of this image with the anatomical high-resolution pre-operative data takes place. The 9 registration parameters are stored as starting position for further linear registration steps. The registered image will be the reference for further steps of the chain.
After the opening of the skull further intra-operative images are acquired. These images are also corrected with respect to possible intensity non-uniformities and registered with the first intra-operative data using the stored parameter set. Before the next step, which is the non-linear registration, will be executed, a linear intensity adjustment might be made to gain the same intensity distribution in both input images. The resulting displacement field of the non-linear registration process is applied to the pre-operative fMR data. In the last step the deformed fMR data will be overlaid to the linear registered open-skull data set and later sent to a presentation device.
The fluid based non-linear registration method produces best results if its input images are originating from the same scanner. That is the reason why the first (closed skull) intra-operative image has been acquired and used as a reference image for further processing and not the pre-operative high-resolution data which is usually not acquired with an Open-MR scanner.
The current processing time for this chain is about 4 h (Intel Pentium III, single processor). In order not to delay the progress of the surgical intervention too much, a maximum processing time of approx. 10 min is acceptable. The time-consuming registration steps are readily parallelisable on shared or distributed memory high performance computing platforms. Because Open MR scanners are only available in some neurosurgical centres the use of intra-operative 3D structural information for guiding the deformation of the functional information will be studied. It is planned to compare the performance of this processing chain computed on single processor machines (for off-theatre validation), local workstation clusters, and remote high performance computing platforms. Grid technology allows switching between environments easily.

Elements of the Image Processing Chain and their Current Status:
The currently available image processing chain focuses on the registration of low-resolution Open-MR data to low- and high-resolution MR data:

1. Data transfer and conversion: Not implemented yet. Data is expected to be in the machine-independent Vista format.
2. Correction of intensity non-uniformities: This step is very time consuming due to the solution of a huge linear system. To improve speed the linear system is only coarsely solved within a multi-resolution framework, which is sufficient to estimate a proper bias field. Further speed improvement is obtained during operating on an enlarged head mask that will be extracted from the given dataset. The biggest improvement of performance was gained by computing the solution of the linear system in parallel. The algorithm was enhanced to work in a multiprocessor shared memory environment with nearly linear speedup. This was realized by an overlap of iterations. So the next iteration is starting suddenly as soon as all necessary data has been computed, although the current iteration has not finished yet. This is the reason why the maximum number of usable processors is bounded by the number of slices in the MR image.

3. Linear registration: The registration of the low resolution intra-operative dataset to a high resolution pre-operative image is realized by maximizing their normalized mutual information (NMI) resp. cross correlation (CC). NMI is used when data from different scanners is going to be registered and CC is used for datasets originating from the same scanner. To achieve a fast convergence, the down-hill simplex optimization algorithm is used, which performs well and does not require any gradient information. Computation time of the gradient of the NMI cost function, for example, would be 9 times larger than the evaluation time of the cost function itself. To improve speed of the registration, a parallel evaluated speculative down-hill simplex was developed who converges exactly as the original method but twice as fast. A further nearly linear speedup was gained by evaluating the cost function (NMI and CC) in parallel by partitioning the data into blocks and assigning each block to a single processor. The results of the processed sub volumes will be cumulated to obtain the cost function value. All the above mentioned in parallel computed parts of the linear registration can be used in a shared and distributed memory environment.
4. Currently, the success of the linear registration depends on the initial orientation of both images because it can not be assured that during optimization the global optimum is found. A strategy needs to be developed to increase the probability of convergence to the global optimum.
5. Intensity adjustment of two scans: The result of the non-linear registration step of the processing chain depends, among other things, on the similarity of the intensities of both images. Since this step is not very time consuming, it can be performed in serial.
6. Non-linear registration: To obtain a deformation field that can be applied to a fMRI image a non-linear registration is required. In this chain a method based on fluid mechanics is applied. The time consuming part here is once more the solving of a huge linear system. To speed this up and to avoid local minima the system is solved using a multi-resolution approach. A further speedup is achieved by solving the system in parallel in a shared memory environment. Here each processor operates on a single "slice" of the system.
7. Application of a deformation field to fMRI data: The deformation field obtained by non-linear registration will be assigned to a fMRI dataset by shifting each voxel of the dataset by its corresponding vector from the displacement field.
8. Overlay of deformed fMRI data with intra-operative dataset: In this step the deformed pre-operative functional dataset that was initially aligned with the pre-operative anatomical MR scan will be overlaid to the registered intra-operative image to show regions of activation with respect to the brain shift.
9. Conversion and transfer to presentation device: Not implemented yet. Data will be output in Vista format.

Status Summary:

Due to the requirement of processing the whole chain (including image acquisition) in less than 10 minutes, large computation power is needed. This will be usually provided through a local cluster with multiple processors. Those clusters are quite expensive, need maintenance and are often under-worked. Instead of a cluster the grid could be used to execute the chain on a currently available computation centre. For grid interaction, a cheap local terminal is fully sufficient.

At the moment the image processing chain takes about 25 minutes on a 2-processor AMD AthlonMP 1.5 GHz machine to be executed, we are very optimistic to decrease computation time to the required 10 minutes with an increasing number of available processors provided by the grid.

3.4.3 Sub-task 4.3: Cranial Radio-surgery Simulation

Application Scenario:
Gamma Knife® Radiosurgery is a non-invasive medical procedure using beams of ionising photons from 201 ^{60}Co sources to treat intra-cranial lesions. The Gamma Knife® unit comes with a treatment planning system, Gamma Plan® that uses an approximate description of the photon interactions within the head of the patient to calculate the energy dose deposited by these photons in the region of the tumour. There is significant benefit to be obtained from improving the fidelity of these calculations, particularly in cases where photons traverse regions of widely differing electron densities (e.g. soft tissue and bone). Monte Carlo calculations provide such a model and can be used initially to complement Gamma Plan®, with the potential to eventually supersede it in the event of short enough calculation times. The goal of this sub-task is to adapt a Monte Carlo code (RAPT/EGS4) written for conventional radiotherapy to that of stereotactic radiosurgery, for calculation of the energy dose delivered to the brain from a Gamma Knife® treatment unit.

Workflow:
This application uses the code RAPT, to extend the Monte Carlo kernel EGS4 with the aim of simulating the energy dose delivered by the Gamma Knife® to the region of the brain tumour. As input, RAPT requires: patient geometry, types of material contained in the geometry, isocentre of the photon beams, beam intensity and energy profiles, number of beams and their angles, and number of photons per beam to guarantee convergence and effective simulation of Gamma Knife® outcome. This is entered as text data and saved as input files to be used by the RAPT software, which converts the data to EGS4 compatible format and initiates modelling of millions of photon paths to calculate volume energy deposition within the patient head. This data is output as text files that are then visualized with routines written in MatLab. The output files contain: 3D patient mesh geometry, accumulated energy dose at each mesh (voxel) element, and the flux statistical uncertainty at each voxel.

Grid Enabling:
The RAPT software has been run and tested under LINUX and Windows NT. For the purposes of the alpha release, Grid-enabled RAPT has been executed remotely using a Windows NT client. The example problems so far attempted have relied upon phantom data – no attempt has been made to import detailed head geometry from CT or MRI. It should be noted that this is consistent with standard clinical practice, since GammaPlan uses a geometry based on a parameterised ‘bubble head’ rather than diagnostic imaging data. Data for the bubble head is derived from 20 measurements of the patient’s skull with respect to a stereotactic frame. The lack of complex head geometry data minimises data transfer requirements so that at most approximately 1MB of data is transferred to and from the Grid for a single simulation run. This is likely to increase to 100MB in the case of refined patient geometry.
An X509 compliant PKI security infrastructure is under incremental construction. The data controller at the hospital has been identified and candidates for the Registration Authority are under discussion. The hospital firewall is very tightly monitored and it is not possible to ‘push’ data from an external source into the NHS network. However, the ability to ‘pull’ the data from within the bounds of the hospital firewall makes it possible (in principle) to run remote Grid jobs from the hospital environment. Data security is paramount and utilises 128-bit encryption. The fact that a patient could be identified from the voxel data set does not so far seem to require special security means beyond that provided by the GEMSS system.
The QoS model lets the user specify the time of job completion and uses an estimate of job run time based on a benchmark problem. Timely solution is critical in this clinical application,

and it is possible to poll the Grid and determine the state of the simulation while it is running. Currently a simple telephone business model has been installed that calculates cost as a function of CPU time used.

Sample Problems and Test Cases:
Attenuation of a single monochromatic photon beam in a cuboid of water has been simulated using RAPT and the results verified against published data (linear attenuation coefficient etc.). Our benchmark case is a spherical water phantom target with isocentre at its centre. The results of the RAPT simulation compare favourably with the dose distribution computed by GammaPlan for the same target. Additionally the limitation of the GammaPlan physics model has been noted by observing differences of dose distribution at a linear attenuation coefficient boundary.

Status Summary:
The achievements to date are:

- RAPT adapted to model the Gamma Knife (beam angles, 201 beams, 1 mm resolution).
- Incremental fixing of instabilities in the RAPT code.
- Creation of visualization tools. RAPT output is visualized using MatLab.
- Successful execution of RAPT on two different platforms (Linux and WindowsNT)
- Successful execution of RAPT on the Grid (GRIA) using a Windows NT client.
- Validation of RAPT using two test cases (water block; spherical phantom).
- Integration of GEMSS-compliant features into RAPT.
- Limited security model
- Business model
- QoS model

Next Steps:
The goals for the project over the next year include:

- Achievement of 0.1 mm resolution.
- Implementation of a visual interface to ease modelling of realistic treatment plans and visualisation of the results.
- Further comparisons of RAPT with GammaPlan. (eg. improved modelling of head geometry (Bubble head); development of methods to facilitate these comparisons)
- Refined modelling of the Gamma Knife:
    - Modelling of different collimators - the Gamma Knife® uses collimator sizes of 4, 8, 14, and 18 mm.
    - Modelling the air gap between collimators and the patient in RAPT. (RAPT accepts Cartesian geometry, while the Gamma Knife® unit has spherical symmetry)
    - Modelling the Co-60 sources and collimators of the Gamma Knife® to improve description of the input beams.
- Implementation of a full security model

3.4.4 Sub-task 4.4: Inhaled Drug Delivery Simulation

Application Scenario
This application is a simulation of inhaled drug delivery (acronym 'COPHIT' – '**C**omputer **O**ptimised **P**ulmonary delivery in **H**umans of **I**nhaled **T**herapies'). Optimisation of delivery

of medication to the lungs is achieved through a suite of tools that allow the whole of the inhaled drug delivery process to be modelled. In order to do this, knowledge of the geometry of the inhaler and the airways is required. The physical characteristics of the medication must also be specified, because this will determine how it is entrained in the air flow and how it is deposited on the airway walls. The flow rate of air in the respiratory system varies over the course of a breath and therefore the timing of the delivery of medication with respect to this flow profile must be accounted for because it will affect the distribution of deposition within the lungs. The software allows drug designers or manufacturers of inhalation devices to experiment at the design stage with device geometry, delivery timing, or physical formulation of the medication, in order to maximise the dose to the desired region of the respiratory tract. A final component to COPHIT is a pharmacokinetic model, in which the computed deposition of drug within the airways can be used to predict how the plasma concentration of a drug changes over time after inhalation.

Workflow

There are several stages in the simulation process. The pre-processing stages require user interaction and are not computationally demanding - these steps are performed on the local client. Device geometry can be obtained from Computer Aided Design data whereas anatomical geometry is derived from a bitmap image stack of a segmented CT scan set, from which a tetrahedral mesh is generated. Inlets and outlets need to be identified on this mesh. The material properties of the gas or gas mixture are specified. Compartment models representing the compliance of the peripheral airways can also be added to the model at this point. The user specifies the properties and initial conditions for all compartments, following which the model information is assembled into a 'DEF' file which is the input for the CFX computational fluid dynamics solver. Auxiliary text files specify the compartment properties and input flow and composition profile. These files are encrypted and sent to the Grid server, where the computationally intensive fluids solution step is performed, and the results files are retrieved later. Visualisation of the flow pattern is performed on the client. Finally, the drug deposition results can be input into a pharmacokinetic model, also based on the client computer.

Grid enabling

The COPHIT software has been tested using Windows and Silicon Graphics IRIX operating systems. Grid enabling of the application yields particular benefits for the fluids solution step and has been tested on the GRIA Grid. Patient data is stripped from the lung mesh before transmission to the Grid server. In addition it does not seem plausible that the lung geometry could identify a patient in the same way as, for example, in ST4.1 (Maxillo-facial surgery simulation) and the standard GEMSS security precautions are considered adequate. It is possible that the input data could contain information that is commercially sensitive, such as the geometry for a new design of inhaler, and therefore 128-bit encryption is used to ensure the security of the data during transmission. Because this software is used as a research tool, solution is not time-critical, unlike some of the more clinically-orientated applications. A typical job might take of the order of a week to run. It is anticipated that Grid resources would be used at periods of low demand in order to reduce the cost of running the job, which is currently based on a telephone service model.

Test case

The test case models the tracheo-bronchial tree, from the trachea down to the 4^{th} generation airways. A time varying input flow rate is specified at the top of the trachea section. The input medium is specified as a mixture of air and water droplets of diameter 1μm to represent the medicament. The outlets of the CFD model are coupled to elastic compartments which model the compliance of the peripheral airways. The results show how the medicament is deposited in the 3D section of the model and also how it is distributed to the peripheral airways. This test case has been successfully run by GEMSS partners ASD, IDAC, and IT innovation.

Status Summary

Current Status of the Alpha Release

- A COPHIT demonstration disc including Powerpoint and html-based presentations has been developed as a means of introducing COPHIT to potential users. It shows the steps in the solution process and demonstrates the potential of COPHIT for solving inhaled drug delivery problems.
- The software has been partitioned into local-client and Grid-server sections. The mesh generation pre-processing stage has been dramatically improved to reduce the memory overhead, allowing it to be performed on the local client.
- The portability and ease of installation of the application has been improved.
- MATLAB was originally required on the Grid server during the CFD solution stage. This is no longer necessary because the code has been re-written using FORTRAN. Thus a MATLAB license is no longer required for Grid servers. However, MATLAB is still necessary locally for pre- and post- processing purposes.
- Proprietary routines obtained from 'Numerical Recipes' have been replaced.
- An example problem for Grid testing has been developed and successfully tested on individual personal computing platforms and the GRIA Grid.

Future Goals

- Improve the running of Grid jobs, including:
    - ability to terminate jobs and re-start jobs
    - provide an estimate of solution time
    - security measures to check incoming files for malicious code
    - compile and test code for LINUX operating system
- Improve the user interface for the pre-processing stages. COPHIT maintains its origins as a research tool which allows for maximum flexibility. However, the very large number of options available when setting up CFX problems can be distressing to non-expert end-users. This can be corrected by development of a user interface that offers the user a choice between a number of pre-generated meshes, material property sets and boundary conditions. Appropriate results plots could be automatically generated on the Grid server and sent back to the user. We are currently investigating the use of EASA as a means of creating such a user interface.

3.4.5 Sub-task 4.5: Cardiovascular System Simulation

Application Scenario

Modelling of blood flow in arteries requires the use of sophisticated three-dimensional computational fluid dynamics (CFD) software. Using such software, flow through isolated vessel sections can be simulated to provide insight into pathologies of the heart and vasculature. However, even if we are only interested in the flow in a particular section, the properties of the whole of the vasculature should be taken into account, because effects like peripheral vasoconstriction can have a profound impact on flow anywhere within the cardiovascular system. Unfortunately, as more of the vasculature is included within the simulation, the number of vessels increases exponentially and the computational problem becomes intractable. Thus it is necessary to reduce the magnitude of the problem by encapsulating the properties of the peripheral circulation in a compartment model. In this way the dependence of the flow on the peripheral circulation can be investigated with much lower computational requirements. This application aims to couple the full three-dimensional CFD model of a vessel section to a terminating compartment and builds on principles developed in ST4.4 The cardiovascular simulation will develop the compartment properties and its method of specification.

Workflow

The workflow of this application is similar to that of COPHIT described previously (ST4.4). In the pre-processing stages, a tetrahedral model mesh of the vasculature is generated from a segmented CT or MRI scan set, and inlet and outlet boundary conditions are specified. The properties of the compartment model coupled to the 3D CFD model are also specified. In this case the CFD model is of a vessel segment rather than airways, but the mathematical and physical principles underlying the problem are very similar to that of COPHIT. The pre-processing stages are performed on the local client, and the input files for the CFD and compartment solvers are exported to the Grid for solution. Following satisfactory solution, the results are returned to the client for flow visualisation.

Grid enabling

The cardiovascular application uses the same Grid-enabling technology as COPHIT (ST4.4). This includes anonymisation of all patient data and 128-bit encryption as part of the GEMSS standard security precautions. The fluids solution step derives great benefit from the Grid, but even so, a typical job might take of the order of a week to run. It is anticipated that Grid resources would be used at periods of low demand in order to reduce cost.

Test case

The test case utilises an idealised bifurcation model, with a tube of diameter 10mm branching to form two tubes of diameter 8mm. Blood has been modelled as a Newtonian fluid of viscosity $4*10^{-3}kgm^{-1}s^{-1}$. An input flow rate that varies with time has been specified. Two compartments have been added to the outlets of the 3D model, representing peripheral resistance and compliance. The two peripheral resistances are different in magnitude, and the results show how this affects the flow pattern, with greater flow evident in the branch of low resistance.

Status Summary

Current Status of the Alpha Release

- A new compartmental model for representing the peripheral circulatory system has been developed.
- The software has been partitioned into local-client and Grid-server sections. The mesh generation pre-processing stage has been dramatically improved to reduce the memory overhead, allowing it to be performed on the local client.
- MATLAB is used locally for pre- processing purposes.
- A simple method for installing the FORTRAN code for the compartment model on the Grid server has been implemented.
- An example problem for Grid testing has been developed and successfully run on individual personal computing platforms and the Grid.

Future Goals

- Improve the running of Grid jobs, including:
    - ability to terminate jobs and re-start jobs
    - provide an estimate of solution time
    - security measures to check incoming files for malicious code

- compile and test code for LINUX operating system

- Develop a user-friendly method of specifying the compartment properties in the most appropriate way – such as compartment transfer function, frequency response, or an electrical circuit analogue of the system.

3.4.6 Sub-task 4.6: Advanced Image Reconstruction

Application Scenario:
Tumour diagnosis and monitoring of metabolism are the main tasks of in vivo diagnosis in nuclear medicine by visualisation of distribution of radioactive tracer in the human body. Although SPECT 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 ionising 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. A large variety of reconstruction algorithms exists, many of them based on standard filtered back-projection (FBP). This method finds extensive use in current clinical practice but it is only applicable to single slices. The modern iterative algorithms available within GEMSS offer benefits because they encompass technical and physical constraints of the imaging process and are easily extendable to 3D, but this comes at the expense of high computational effort. The Grid is well suited to the task in addition to which an implementation of image reconstruction as a Grid service could also bring access to highly sophisticated image processing resources for better diagnosis.

Workflow:
Diagnostic SPECT images are reconstructed from projection data, (ie. the sum of emitted photons along a linear manifold). In practice this manifold is a cone-like sub-volume of the object, from which photons are recorded by surrounding detectors. A rectangular detector array is rotated around the patient and a series of projection data is acquired. Iterative reconstruction repeatedly modifies a postulated image matrix through comparison of pseudo-projection and measured projection data. By this method a succession of intermediate images is generated until a convergence criterion is fulfilled. The weighted contribution of each voxel of the image to a specific projection value permits accurate modelling of collimator geometry and photon scatter, and all weights for all projection values define the system matrix which characterises the system response. If N is the length of the image cube, then the number of calculations is of order N^6 for full 3D reconstruction compared to order N^4 with slice reconstruction.

Grid-enabling:
The purpose of this application is to provide high level image reconstruction in emission tomography with transparent access to high performance computing resources via the Grid, including specific security measures and billing models. The reconstruction service is implemented as a parallel reconstruction kernel embedded in a service oriented environment.

Reconstruction Kernel:
A state of the art algorithm - the ML-EM algorithm of Shepp and Vardi – has been implemented. This algorithm is based on a stochastic model of Poisson distributed generation and detection of photons. Well known stability and robust convergence criteria make this the algorithm of choice. The algorithm was implemented in ANSI-C employing a hybrid

parallelization paradigm optimised for SMP-clusters. Symmetric shared memory parallelization was realised using OpenMP directives. Inter-node communication was implemented using MPI from the mpich library.

Applications Provider's View:
The reconstruction kernel was embedded in a service oriented environment. The code is exposed in our system as a Web Service. Services may be accessed via a browser-based graphical user interface either in direct mode, in which case a predefined reconstruction service is accessed, or, in managed mode, where a service manager selects a service on behalf of the user. In this preliminary implementation services are predefined and only rudimentary selection strategies are implemented, e.g. low workload at the machine. Patient data are stripped on the client side and kept private, thus identification of the patient by projection data is impossible on the server machine.

Example Problem / Test Case:
The test case does not utilise patient data but uses a phantom to allow for quantitative evaluation of the accuracy of the 3D method. Projection data from a Jaszczak phantom was reconstructed and compared to filtered back-projection (FBP). The phantom contained 600MBq Tc99m and projection data was acquired on a GE Hawkeye SPECT scanner with a 128x128 projection matrix, 3 degrees rotational increment, 30 s acquisition time per step. Images were reconstructed using FBP with a Hanning window and compared with fully 3D ML-EM reconstruction using an optimised system matrix with the alpha release software of the GEMSS project. From this phantom data the enhancement of contrast in both sections of the phantom, i.e. cold spheres and cold rods, could be shown.

Status Summary:
Status of the Alpha Release
1. development of a parallel reconstruction kernel using a hybrid parallelization strategy
2. implementation of the kernel on SMP clusters
3. integration of the kernel in a preliminary service oriented environment
4. development of a service oriented environment based on Web Services
5. development and implementation of a GUI for SPECT reconstruction
6. validation of fully 3D reconstruction with a clinical standard test phantom
7. demonstration of the feasibility of fully 3D reconstruction as a Web Service

Future Goals
1. development of an accelerated reconstruction algorithm using ordered subsets
2. calibration of the system-matrix to standard collimator types and (most importantly) clinical isotopes
3. connection of camera systems based on the DICOM standard
4. validation of the reconstruction service in a clinical environment
5. implementation of the GEMSS authentication and billing procedures implementation of the service in the OGSA in collaboration with partner ISS.

3.5 Workpackage 5: Exploitation, Information Dissemination and Clustering

This workpackage is executing tasks related to exploitation and information dissemination. In particular it will plan and co-ordinate information dissemination activities among the partners, as well as produce a post-project exploitation roadmap. A second task will support project clustering activities along two axes: generic GRID technologies and technologies for applications in the health care domain.

3.5.1 Exploitation Planning

The GEMSS consortium has broad exploitation and use intentions for the project output, covering non-commercial and commercial exploitation. The details of the GEMSS exploitation plan, documented in a separate internal deliverable – Exploitation Planning Report, D5.1a, are confidential. D5.1a was prepared under Work Package 5 Exploitation Planning and Information Dissemination of the project, and complements the material in the eTip (http://etip.cordis.lu) containing the Technology Implementation Plan. The D5.1a report sets out to describe planning of the exploitation of the results of the project. Particular attention will be given to the markets in which GEMSS exploitation can take place, the opportunities within these markets, the barriers which would need to be overcome, and the planning required to exploit these opportunities. The dissemination activities necessary to complement the exploitation activities are described separately in the Dissemination and Use Plan, D5.2. The exploitation planning report firstly presents a very general description of the exploitation potential of the project. These general notes will be followed by the description of domain specific activities, and the description of individual partners' specific plans, and how they can be integrated for maximum benefits.

3.5.2 Dissemination Activities

GEMSS project partners observe current developments in the area of GEMSS activities through the attendance of conferences and events to ensure that the public will be informed about the work of GEMSS and that GEMSS will be aware of other projects´ activities. Relevant conferences and meetings including events organized by the European Commission are listed in table 1 below (period 01-SEP-02 to 31-AUG-03).

Name	**Where**	**Start**	**End**
Summer School: The lungs: function diagnosis and treatment.	Oxford (UK)	06-JUL-03	11-JUL-03
IST Grid Concertation Meeting	Brussels (Belgium)	18-JUN-03	19-JUN-03
Conference “Urgence aux Urgences : l’informatique, une solution?”	Namur (Belgium)	05-JUN-03	05-JUN-03
Training in Medical Data Management	Brussels (Belgium)	27-MAY-03	28-MAY-03
Short training in Medical Data Management	Brussels (Belgium)	11-MAY-03	11-MAY-03
Certificat Inter-Universitaire de Traitement de l’Information Hospitalière (C.I.T.I.H.)	Louvain-la-Neuve (Belgium)	09-MAY-03	09-MAY-03
Seminar « Droit et éthique de la recherche médicale, Originalité de la législation française, limites et questions éthiques »	Toulouse (France)	06-MAY-03	07-MAY-03
Meeting w. Health Telematics Commission	Brussels	24-APR-03	24-APR-03
International Parallel and Distributed Processing Symposium - IPDPS	Nice (France)	22-APR-03	26-APR-03
Conference of the « Association Belge des Directeurs d’hôpitaux »	Mont-Godinne (Belgium)	03-APR-03	03-APR-03
Conference « La Judiciarisation de la medicine »	Paris, Université Necker	10-MAR-03	10-MAR-03

Global Grid Forum, GGF7	Tokyo (Japan)	04-MAR-03	07-MAR-03
IDAC Ireland Medical Device Seminar	Galway (Ireland)	23-JAN-03	23-JAN-03
GEMSS Project Presentation	Nice (France)	21-JAN-03	21-JAN-03
1st European HEALTHGRID Conference	Lyon (France)	16-JAN-03	17-JAN-03
Seminary JURITIC	Namur (Belgium)	12-DEC-02	12-DEC-02
2nd Cracow Grid Workshop	Cracow (Poland)	11-DEC-02	14-DEC-02
Conference « Informatique et Télématique des Soins de Santé »	Brussels (Belgium)	05-DEC-02	07-DEC-02
Seminar - Clinique Sainte Elisabeth	Namur (Belgium)	26-NOV-03	26-NOV-03
IST2002 - Partnerships for the Future	Copenhagen (Denmark)	04-NOV-02	06-NOV-02
iGrid 2002 - The International Virtual Laboratory	Amsterdam (The Netherlands)	23-SEP-02	26-SEP-02
UNICORE Workshop	Bonn (Germany)	27-NOV-02	28-NOV-02
Conference on Genetics and Law	London (U.K.)	19-OCT-02	20-OCT-02
CAE Consortium Meeting	Friedrichshafen (Germany)	12-OCT-02	13-OCT-02
CAE Users Conference	Friedrichshafen (Germany)	09-OCT-02	11-OCT-02
EC Grid workshop	Brussels (Belgium)	03-OCT-02	04-OCT-02
CFX User Conference	Strasbourg (France)	17-SEP-02	17-SEP-02
HEALTHGRID meeting	Brussels (Belgium)	20-SEP-02	20-SEP-02
GEMSS Kick-Off	Brussels (Belgium)	19-SEP-02	19-SEP-02

- **Contacts with industry**
  GEMSS has links to Industry through contacts established in other projects and existing contacts especially by the industrial Consortium partners who also disseminate information about GEMSS related work to their (industrial) user communities.

- **Collaboration with other projects**
  GEMSS participates in the HEALTHGRID cluster formed by EU projects in the e-Health / Grid area.

  GEMSS has established a link with the GRIDSTART project. Partner IT Innovation is a contractor of the GRIDSTART IST project. NEC representation of GEMSS was offered to a representative of GRIDSTART.

- **Information sources**

  - **Websites**

    The public GEMSS website http://www.gemss.de was designed according to the EC DG Information Society guidelines version 1.1. It provides general information about the project. The index page offers the following links: Home, Project Details, The Consortium, Reports and Presentations.

Further material on the project is contained within a secure website, with partner access only. It provides publicity material which can be used by individual partners for their own dissemination activities.

- **Printings**

  GEMSS has so far released:

  - a Project Handbook available to IST projects,
  - a general brochure detailing the Project,
  - a general Project poster,
  - a publication "Medical Simulation Services via the Grid", proceedings of the first HEALTHGRID conference in Lyon, 16.-17.1.2003,
  - other media articles (partner presentation, application brochure, CD) prepared on a by-request basis.

3.5.3 Clustering

GEMSS participates in the HEALTHGRID cluster formed by EU projects in the e-Health / Grid area. Partner NEC is an official member of the HEALTHGRID Management Board.

GEMSS has established a link with the GRIDSTART project. Partner IT Innovation is a contractor of the GRIDSTART IST project. NEC representation of GEMSS was offered to a representative of GRIDSTART.

USFD links have enabled GEMSS to raise its profile within the White Rose Grid (UK). Furthermore, GEMSS is directly cooperating with the EC project Gria and the UK EPSRC project MyGrid. GEMSS also cooperates with the GridLab project through the Triana group at the University of Cardiff.

3.6 Workpackage 6: Project Management

This workpackage co-ordinates day-to-day running of the project and maintains all correspondence with the European Commission. The main goal of WP6 is to ensure the project's success by monitoring and reporting progress against goals and time scales, by establishing processes to ensure quality in project work and by anticipating and managing risk and change to the project.

3.6.1 Project Communication

The communication strategy aims to keep all the Partners fully informed about the Project status, the planning and all other issues which are important to the Partners in order to obtain maximum transparency for all involved and to increase the synergy of the co-operation.

Interactive management meetings and technical meetings take an important role in the communication strategy:

- All information (like minutes of meetings, visit reports, tasks reports, relevant publications etc.) is communicated to the Project Co-ordinator, who is responsible for also channelling this information to the Partners, where appropriate.
- BASIC SUPPORT FOR COLLABORATIVE WORK (BSCW): A BSCW server has been set-up as a transparent shared work-space at the co-ordinators site. The server complements the web-server and contains all project internal documents. In addition, it serves as a common repository for documents used in preparation of deliverables. Currently 41 users from all partners are registered.
- Email reflectors have been set-up at NEC for the following (sub)groups:

All partners:		gemss@gemss.de
PMB:		pmb@gemss.de
WP1:		wp1@gemss.de
	ST1.1:	st11@gemss.de
	ST1.2:	st12@gemss.de
	ST1.3:	st13@gemss.de
WP2:		wp2@gemss.de
	ST2.1:	st21@gemss.de
	ST2.2:	st22@gemss.de
WP3:		wp3@gemss.de
	ST3.1:	st31@gemss.de
	ST3.2:	st32@gemss.de
	ST3.3:	st33@gemss.de
WP4:		wp4@gemss.de
	ST4.1:	st41@gemss.de
	ST4.2:	st42@gemss.de
	ST4.3:	st43@gemss.de
	ST4.4:	st44@gemss.de
	ST4.5:	st45@gemss.de
	ST4.6:	st46@gemss.de
WP5:		wp5@gemss.de

3.6.2 Management

The project management has acted as contact point for all correspondence between the project and the Commission. Project-internal communication has been simplified by the provision of GEMSS mailing lists for each workpackage or subtask. All correspondence via the GEMSS lists is automatically archived. The Consortium Agreement has been constructed, a modification of the Unified Consortium Agreement, and agreed between all partners. The Agreement will be signed as part of the pending Contract Amendment.

Two full project meetings took place in the reporting period. Project Management Board (PMB) meetings were held during the two-day meetings. The project meetings were used to define and review workplan implementation details and collaborations and also to address exploitation possibilities. The regular meetings were accompanied by several technical sub-group meetings as required by the project.

A cost neutral change of the effort distribution of the Max-Planck-Institute for Cognitive Neuroscience (MPI) was suggested on September 23rd 2002 and accepted by the project officer on October, 3rd 2002.

After the first quarter MPI reported a hiring problem. A recovery plan documented in the 2nd and 3rd Management Reports also communicated to the project officer was approved by all members of the GEMSS Project Management Board.

In the first management report it was announced that AEA Technology is focussing its attention on its core businesses in the Environment and Rail, and is divesting some of its non core activities in order to be able to use the income to invest in its core activities. As a result, the CFX business, within AEA Technology Engineering Software, has been sold to ANSYS, an American company specialised on stress analysis. CFX and EASA business units will be separated into different companies. The restructuring will require a contract amendment. Partners and Project Officer have been informed at the earliest possible stage at the end of the first quarter. The contract amendment is in progress.

The first cost claims of the project corresponding to the reporting period September 2002 – February 2003 were submitted on time on April, 30th 2003.