

Use of Grid technologies for EIT reconstructions

*Marc Molinari, Simon J Cox
University of Southampton
Southampton, SO17 1BJ, UK
m.molinari@soton.ac.uk*

We demonstrate the use of computational Grid technology applied to the area of EIT. In particular the use of databases to re-use existing models and compute resources to run simulations in parallel for the determination of convergence properties is presented.

Introduction

With the ever-increasing size of geometric models and requested detail of accuracy in forward and inverse simulations, the computational demand of EIT reconstruction algorithms can be very high. In particular in the biomedical field, where simulations require non-simple geometrical 3D models with many finite elements, accurate solutions can take a long time to compute and may require large amounts of memory in the reconstruction process.

In addition, every simulation run requires a large number of input specifications such as geometry, mesh files, electrode positions, current injection details, measurement pattern and data, algorithm parameters, convergence results and final reconstructed images.

To be viable in extensive research work on a day-to-day basis, any software used for EIT reconstructions will need to address these two issues of computational requirements and data storage, search & retrieval functionality.

We address these two issues and present techniques on how to implement currently available technology to cope more easily with the above demands. We present a possibility of how to make use of high-throughput Grid technology such as Condor (and Globus) computational resource pools. We show how to distribute the computational load of a full reconstruction onto a pool of networked computers using simple Matlab (or Jython) scripting technology.

The effective storage and subsequent re-use of previous results is our second aim. This relies on the storage of metadata from the reconstruction process and a parametric geometrical model in a database. A best-fit solution can be suggested and pre-selected either automatically or by the user. This technique also allows for better local convergence of the algorithm under the steering control of the experience of a biomedical expert.

Compute & database infrastructure

In many large-scale scientific applications, the compute power required to perform simulations can be huge, especially if results are required immediately or in the near future. In the last decade, there has been a development towards the Computational Grid, which promises to make high performance compute and data storage facilities easily available to desktop end-users. In the UK, the e-Science programme has been instigated in 2001 with the aim of providing middleware, software interfaces and hardware resources to the scientific community. As part of this programme, we have implemented a number of interfaces for Matlab and Jython in the Geodise project (Grid-enabled optimisation and design search, <http://www.geodise.org>) and the GEM project (Grid-enabled electromagnetic optimisation). We aim to make available access to resources on the National Grid Service (NGS) through a Matlab interface which allows Matlab users to harness the power of the compute and data facilities provided by the NGS. Locally at the University of

Southampton, we use a Condor pool of 1500 networked desktop PC's running Windows XP and a further 800 Linux-based processors over the Globus interface for our research. As data storage facility, we are using the Oracle 9i and 10 databases and for file storage a GT2 server. All data access and transmission is protected by X509 certificates which we use for transmission, storage authentication and authorization control.

Parametric model EIT example

EIT simulations are based on many parametric properties. Figure 1 shows some of these which are usually stored in files. Any of these parameters can affect the outcome of a reconstruction. If a very similar example has been computed before, the user might have to explore a large range of existing runs to determine if any of these is closely enough related to be used as starting configuration for the run on the new measurement data. Database technology and compute resources can help to identify these close simulations and run a range of simulations in parallel to obtain a good overview of the convergence properties of the algorithm.

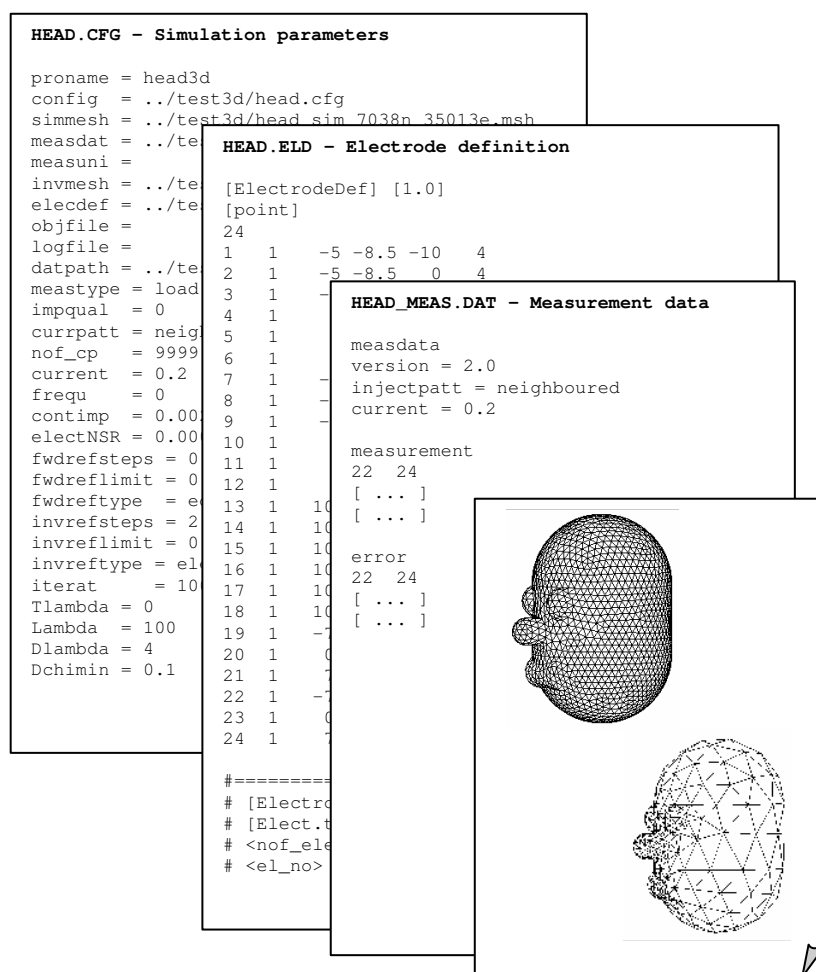


Figure: Metadata and data files needed or produced during a single EIT reconstruction.

The following sample script shows how to search for and submit a suitable range of geometric models for simulation with a compiled executable (e.g. EIDORS-3D or your own reconstruction software) on a Condor cluster. These geometries can be approximated from previous simulations or from simulations of similar patients or the knowledge of certain parametric impedance values. We aim to find a range of solutions for these starting values which then can be evaluated by the medical expert to draw a conclusive opinion on the condition of the tissue distribution.

We aim to re-use results from previous investigations and to obtain a good fit from either a parametric model or a previous model. Geometries and parameterizations can easily be chosen by specifying the range of interest for queries of data stored previously in the database.

```
% find all previous runs in database which match range
models = gd_query( ' model_name==head & eps_skull>100 & eps_skull<140 &
                  skull_thickness>0.80 ' );
for c = 1:length(models)
    % download file to local model directory
    gd_retrieve( models{c}.standard.ID, ['./model_', num2str(c)] );
    % create submission information
    job = condor_job('beginner');
    job.executable = 'reconstruct.exe'
    job.arguments = 'model.geom';
    job.inputdir = ['./model_', num2str(c)];
    jobhandles{c} = condor_submit( job );
end

% wait for results
condor_waitfor( jobhandles );

% display results
display_reconstructions( './model_*' );
```

In this script, the variable `models` contains the number of matching database entries (previously added to the database with `gd_archive`). The relevant files are then downloaded into a local model directory and in the for-loop a job is defined to run the reconstruction based on the model geometry. These jobs are then submitted to a Condor cluster to be computed in parallel.

The results then can be collected and displayed so that a user with medical expertise can quickly identify the results which seem to be most fitting. In the engineering field, this method is similar to a design search and allows users to quickly establish the fitness and stability of a presented solution as there are many local minima the search algorithm could converge to.

Grid technology allows extracting images for a large number of differing models in a time similar to that of a single simulation. In addition, the larger memory requirements of 3D models can be matched more easily by the resources available through the Grid and thus allows the use of finer meshes and more accurate geometric models.

Summary & Future Work

We have addressed the issues of compute and data storage requirements in 3D EIT with the application of Grid technology. Existing and new EIT reconstruction algorithms can be run in a compiled version on a Condor compute cluster in parallel and be submitted and controlled from within the Matlab scripting environment. Geometric models with parameterised shell sizes and tissue impedances can easily be selected from a range of pre-existing models stored in a database.

We plan to integrate a framework which will make this work even more transparent in a day-to-day reconstruction process, suitable for practitioners in EIT. We propose a collaborative project for the application of distributed compute and database technology for EIT practitioners and researchers. If you are interested, please get in touch with the author via email.