Remote Computed Tomography Reconstruction, Simulation and Visualization Service

Overview *Summarise the context that leads to this project.*

The Australian Synchrotron (AS) is a major scientific facility that allows non-destructive high-precision investigation of the internal structure and properties of objects. Light from the synchrotron is directed into beamlines, each equipped with specific scientific instruments dedicated to the conduct of particular types of experiments.

The Imaging and Medical beamline (IMBL) is a major initiative at the Synchrotron and will provide a wide range of advanced X-ray imaging techniques for the biomedical and materials science research communities. About 50% of the available time on the beamline is scheduled to be devoted to 3D Computed Tomography (CT) imaging. The majority of that will be imaging of biomedical samples. It is envisaged that, in the longer term, IMBL will also be used for clinical medical applications, e.g. mammography. IMBL is anticipated to be heavily over-subscribed by users from Australia and overseas right from commencement of operations and therefore high productivity of the beamline is a crucial issue. IMBL is equipped with a very high-spatial-resolution X-ray digital camera with fast readout speed. At the maximum resolution, image data sets for CT reconstruction will have at least 4000 x 4000 x 4000 voxels (128 gigabytes for 12/16 bit pixels). A 3D data set will typically be generated in 15-20 minutes. This National Facility greatly augments the capabilities of the growing tomography community in Australia. Many complementary tomographic techniques (electron, conventional X-ray, optical) will make far better use of the AS through a standardised approach to data processing, storage and visualisation.

High-resolution 3D CT reconstruction, analysis and visualisation typically requires too much memory and compute power to be done on a desktop PC. A super-computing cluster at the Synchrotron (called "MASSIVE") has been funded (members of the present proposal were closely involved in the successful funding proposal to the Victorian Government) and will be purchased later in 2009. MASSIVE will have at least 256 nodes with at least 1024 CPU cores. About 25% of MASSIVE time will be allocated directly to IMBL with the goal to reconstruct high-resolution 3D images at about the same rate as the data acquisition. Time on MASSIVE is also planned for remote CT work, i.e. for CT reconstruction and subsequent 3D data analysis by users who have already completed their experiments at the AS and are then given an opportunity to further process the data acquired in the course of the experiments by accessing the MASSIVE cluster from their home institutions (as opposed to "local" access from within the Synchrotron facilities). Specialized scientific software is required for CT reconstruction at IMBL, as it will be utilizing a number of advanced phase-contrast CT modes (e.g. propagation-based, analyser-based or grating-interferometer-based phase contrast).

Novel applications, particularly in the biomedical domain (low X-ray dose) and material sciences (highly anisotropic samples) will require more specialized reconstruction algorithms, e.g. variants of the so-called limited-angle and local CT reconstruction, tomosynthesis and laminography. CSIRO has a world-leading team developing novel methods for phase-contrast CT image reconstruction. CSIRO has also recently built a working system for delivering remote CT reconstruction and simulation services over the Internet which can be used as a prototype for the current project.

The main focus of this project will be to develop a service for remote 3D CT reconstruction, simulation, analysis and visualisation for the IMBL beamline at the Australian Synchrotron. The framework that will be utilised in this project is more generally applicable, and this will be demonstrated by also implementing these services at another site, expected to be the Clayton Imaging Institute or (if sufficient funding was available for the project) at the ANU CT facility in Canberra.

Users

Identify the research communities and resource providers that this project serves; and the potential number of users. This should include some NCRIS capabilities or other data federating or collaborating research groups, and any institutions that will participate through setting requirements for or steering this project.

Teams of scientists from all over Australia and New Zealand, from universities, hospitals, medical research institutes, CSIRO, ANSTO and other research institutions plan to use and are already using the Synchrotron. The various instruments currently available are attracting considerable interest and requests for use from scientists worldwide. This interest arises from special capabilities available or planned at the AS that are not available elsewhere.

The main relevant area of applications for the technology involved in the present proposal is biomedical imaging in general, and quantitative phase-contrast X-ray imaging and CT in particular. Here, 2D and 3D imaging of small animals and, eventually, human patients will be performed at IMBL, with a particular focus on lung, brain and breast imaging, as well as bone densitometry and the study of subtle changes in bone architecture.

Other areas of the CT research include materials science applications (high-throughput characterization of novel alloys, microelectronics components and nano-composites), geosciences (3D imaging of minerals and oil-bearing rocks), wood science (investigation of effects of climate change on wood structure and properties, etc.), paleontology, forensic studies, food science, etc.

Current Australian users of this imaging modality at overseas synchrotrons (while IMBL at the Australian Synchrotron is being constructed) include:

1. University of Melbourne: understanding osteoporosis through CT imaging of bone.
2. Women's and Children's Hospital Adelaide: rapid, non-invasive assessment of airways in respiratory diseases, cystic fibrosis.
3. Monash University: lung aeration following birth and infant resuscitation, tracking of brain tumour cells tagged with gold particles.
4. Peter MacCallum Cancer Centre and Monash University: treating cancers with synchrotron microbeam radiation therapy.
5. CSIRO: image processing and advanced CT methodology, application to geosciences, biomedical and material sciences, etc.
6. ANU: world-leading research in conventional micro X-ray CT and its applications.

All these and many other users stand to benefit from the proposed remote CT reconstruction service.

Needs

Describe the needs of the research communities or resource providers that this project seeks to address.

It is estimated by the Synchrotron that IMBL will be used about 50% of the time for CT work. The need for a remote CT reconstruction and visualization service at the Synchrotron is driven by the following factors.

Synchrotron experiments are generally quite expensive due to the high cost of the equipment, associated travel costs, etc. Each synchrotron user group is typically allocated a limited time slot for their experiments (typically from 8 hrs to a few days). It is critical for the success of experiments that the preparation and planning stage is as thorough as possible. CT experiments require more synchrotron time compared to conventional 2D imaging. The availability of remote CT services will help the researchers to simulate and test the planned experiments "in silico" in preparation for the real experiments. Our software is capable of simulating CT experiments as performed at synchrotron beamlines or with laboratory-based "conventional" X-ray sources by computing X-ray propagation through 3D models of arbitrary samples presented as 3D spatial distributions (sets of voxels) of the complex X-ray refractive index in the sample, or as 3D distributions of a set of materials with known chemical composition, for an arbitrary set of projection angles (corresponding to different rotational position of the sample during the CT scan). Different statistical noise levels in the data can be also simulated. The simulated X-ray CT data produced in this manner can be used for CT reconstruction just as any "real" experimental data, thus providing a closed-loop "simulation & reconstruction" cycle that can be used for modelling realistic X-ray CT experiments. Additional modules in our software further facilitate the process by allowing the user to modify, visualize and analyze the data at any stage of the CT simulation and reconstruction process.

During the experimental stage, the remote CT reconstruction and visualization services can significantly increase the cooperation between the small number of researchers typically conducting the experiment at the Synchrotron and other members of the research team at different locations in Australia or even overseas. Quick pre-processing, reconstruction and visualization of experimental CT data will allow the researchers to make the necessary adjustments *during* the experimental shift and modify the experimental conditions accordingly. This will significantly increase the productivity of synchrotron-based CT compared with the conventional mode of operation, where the CT reconstruction is typically performed "off-line" *after* the experiments, when the researchers have returned back to their home base.

The 1D (linear cross-sections) and 2D visualization (single plane cross-sections and 2D "movies") will be provided by the facilities built into the main application (these facilities will be available both in the remote modules running at the cluster and in the rich client module running on the user machine). The 3D rendering and visualization will be handled by third-party program(s), most likely by the free program DRISHTI developed at the ANU Vizlab. In the "remote" mode, the 3D visualization will take place at the Synchrotron (using specialized cluster nodes equipped with high-end GPUs) and transmitted to the user machines via the Internet connection VNC-style (i.e. by transmitting the desktop contents). This approach allows us to deploy other 3D visualization tools at the cluster alongside with DRISHTI if desired. VeRSI has developed improved VNC-style tools that provide better security and improved performance compared to standard VNC implementations. These tools can be used for the purpose of remote 3D rendering and visualization as described above. Sharing of 3D visualizations between different sites will be considered, but may be problematic due to the very high network bandwidth that it is likely to demand.

In the post-experiment stage, the access to centralised supercomputing facilities running efficient CT reconstruction and visualization codes will allow for optimal post-processing and analysis of the experimental CT data. At present, different user groups have access to a variety of commercial and "home-grown" CT reconstruction codes. These programs are of variable quality, efficiency and reliability, and typically suffer from data format incompatibilities, lack of particular features and other deficiencies. There is a clear need among synchrotron and laboratory CT user communities for access to well-tested, flexible and efficient CT reconstruction and visualization tools running on high-performance computer hardware and providing CT services based on flexible data formats and data handling mechanisms.

3D CT data sets collected with the use of modern X-ray CCD detectors (which typically contain 4K*4K pixels or more), such as the one installed at the IMBL, are typically several hundred Gigabytes in size or larger. The numerical complexity of CT reconstruction algorithms is generally of the order of N^4, where N^2 is the total number of detector pixels. Therefore, even a well-implemented CT reconstruction algorithm takes many hours or even days for the reconstruction of the 3D structure of a sample from such large datasets if executed on a typical desktop computer with up to 8 CPU cores. On the other hand, the CT reconstruction algorithms can be almost trivially parallelized and the reconstruction time scales down linearly with the number of CPU cores. Therefore, with an optimal implementation a CT reconstruction on a 1024-core cluster can be done approximately 1000 times faster than on a single CPU core, and more than 100 times faster than on 8 cores. The amount of RAM required for these algorithms also scales linearly with the number of cores, so it is necessary to maintain the ratio of between 100MB to 1 GB of RAM per CPU core depending on the detector size and the algorithm implementation. Recently, there has been a lot of activity in the area of GPU implementation of CT reconstruction algorithms. Speed-up factors of between 10 and 100 compared to a typical CPU core have been demonstrated with the use of high-end modern GPUs. Our prototype software already has the ability to perform the CT reconstruction on heterogeneous CPU/GPU clusters taking full advantage of the computational power of GPUs where they are available. CSIRO currently has an active research program in the area of GPU-based image processing and other computing problems, including CT reconstruction. Reducing the CT reconstruction time by a factor of 100 to 1000 will allow the users to perform the reconstruction from the data acquired at IMBL and comparable facilities in realistic times.

Providing a central, shared compute resource for 3D CT reconstruction, simulation, analysis and visualisation is therefore of significant benefit to the user community. However this requires a remote access framework so that these services can be easily utilised from the researcher’s desktop, and support shared visualisations of 3D CT reconstructions by collaborators at different locations.

Services *Describe the result of the project in terms of the service(s) that will be implemented and demonstrated by the project and which could be operated in an ongoing fashion; and the proposed operator of each service.*

We propose to develop a software service including:

1. A system for rapid CT simulation and reconstruction from the data collected at the AS and at the second facility (Clayton Imaging Institute or the ANU CT facility). The computations will be performed at the compute clusters installed at the CT facilities.
2. Tools for remote access to the reconstruction and simulation software via Internet.
3. Tools for remote visualization and collaboration capable of handling large images over Internet via network connections with realistic bandwidths.

The software will be designed as a reliable “production system” (as opposed to research or experimental software) with primarily "unsophisticated" users in mind, with the users isolated from the complexities of the communication protocols and details of cluster computing. Instead the users will be presented with a set of standard GUI interfaces that will allow them to perform CT reconstruction in all the modes provided by the CT facility by following the prompts in the GUI. We will provide extensive Help system both on-line and as part of the rich client component, which will include step-by-step tutorials describing the typical reconstruction procedures. The software will be flexible enough to cover the needs of the majority of users in a standardized and efficient fashion.

The "server-side" software components will be running on the super-computing cluster at the Synchrotron. The "rich client" software components will be running on remote user machines. The communications between the server and user components will be based on the standard SOAP web services protocol. Although main components of our software are / will be written in the ISO standard C++, we are considering the multi-platform deployment to be outside the scope of the current proposal. Both the client and the server components will be running under Windows (XP/Vista/Windows 7 in the case of the client components and Windows Cluster Server 2003 / Windows HPC Server 2008 in the case of the cluster components). We understand that a multi-platform implementation including at least MacOS and Linux could be desirable, but it is not feasible to include such implementation into the scope of the present project given its time (2 years) and funding limitations. We will make provisions for possible future porting of our software to different platforms. It is important to note that according to the funded proposal, the MASSIVE cluster will support both Linux and Windows OS. Given our intention to target primarily the "unsophisticated" CT users interested in routine reconstruction, we believe that the Windows-only implementation of the remote CT services will fit the purpose well.

As the main components of our software (including all the remote interface components and specifications) will be open-source, extension and modification of the system by various groups in the future will be possible. The software is modular and can support integration of different reconstruction software and other components.

The proposed CT service will be developed by CSIRO in collaboration with VeRSI and the AS, and the AS will be the long-term service provider. The AS will assist this development with both conventional and computing resources.

We propose to create a *second host site* for provision of remote CT services as part of the present proposal either at the ANU CT facility or at the Clayton-based Imaging Institute, a joint venture between Monash University and CSIRO, which will have several laboratory X-Ray CT instruments. The ANU CT facility may initially install only a small-capacity Windows cluster, with the option to expand it if the user demand for the remote CT services and the operational conditions at the facility will make it possible. Comparison of results with the ANU reconstruction code, currently run on APAC, may assist with improving efficiencies in both systems. The Clayton Imaging Institute will have its own Windows-based computer cluster suitable for hosting the remote CT services. The Institute will also likely have access to the MASSIVE-2 cluster with the possibility to install the remote CT services on that cluster as well. Subsequent deployments at other CT facilities (including AMMRF nodes) may be possible following the model described in the present proposal.

CSIRO has been supporting the development of the CT reconstruction software that will be used for the purpose of the present proposal for more than 8 years. CSIRO has approved plans in place for continuing support of this software in the future. CSIRO will be open to collaboration with other organizations that may decide to participate in the use and development of this software. Routine day-to-day support of the users of remote CT reconstruction services will be the responsibility of the organisations running the services. This approach has been discussed with the proposed deployment sites and preliminary agreements have been reached. As mentioned above, extensive Help system and detailed tutorials will be prepared in the course of this project.

We will provide free access to the developed software for registered users of the CT reconstruction facilities for a period of at least 3 or 6 months after completion of their data collection beamtime. This free access period is planned to coincide with the period during which the Synchrotron or other CT facility will allow the users to access their compute cluster installations (such as MASSIVE) after the completion of the CT experiments at the facility. We will make available through an open source licence the source code of the estimated 90% or more of the pre-existing code base developed at CSIRO that will be used in the proposed implementation of the remote CT reconstruction services. In particular, we will open the source of all components related to the remote access and cluster control. The pre-existing components that cannot be made open-source are primarily those involving patented IP related to X-ray phase retrieval (part of the specialized CT reconstruction mode) owned by CSIRO or third parties. All the source code developed in the course of this project will be made open-source.

Although we propose to initially develop this remote image reconstruction and visualization service for the needs of IMBL users, it is expected that similar services will be useful for other synchrotron beamlines, as well as for other major Australian research facilities (see Broader Adoption below).

NeAT Characteristics

eResearch effect

What changes in behaviour and activity are expected from the project that will demonstrate the broader adoption of eResearch practice?

The availability of the remote CT reconstruction and visualization service at the Synchrotron will provide the following benefits and changes in user behaviour.

1. Increased productivity – remote access to the CT simulation and reconstruction tools will allow users to better plan their synchrotron experiments. It has been estimated that the cost of operating of IMBL will be around $50,000 for a typical 48 hour session. A poor set up can often result in degraded or unusable data and ruined experimental subjects. If improved experimental preparation using remote CT service increases productivity only by 25%, it will equate to a saving worth $586,000 per year.
2. Increased cooperation on experiments – the scientists will be able to use the remote collaboration tools during experiments at the AS in a typical situation where at least some team members are located far from the Synchrotron at the time of the experiment.
3. Increased quality of scientific outcomes – the proposed service will be used after the experiments have been completed for post-processing of the images collected during the experiments via remote access to the Synchrotron computer cluster. This will also bring the CT user community closer together and help increase inter-institutional collaboration by means of adapting common toolsets, common data standards and management practices.

Broader adoption

Which additional communities, resource providers or organisations would also be expected to benefit from the provision of the same or similar services should the project succeed?

Apart from the users of IMBL at the AS, the following groups are expected to benefit from the outcomes of this project.

1. Users of other synchrotron beamlines which are planning to implement CT modes of data acquisition and reconstruction in the future, namely Microspectroscopy and Protein Crystallography beamlines. When a CT or similar modes are enabled on these beamlines, they will be able to utilize the remote CT services infrastructure that will be developed and deployed at the AS as a result of the present proposal.
2. Multiple (more than 10) laboratory-CT-equipped research facilities across Australia (e.g. CSIRO Materials Sciences and Engineering, La Trobe University, ANU, etc., many of these being members of Australian Microscopy & Microanalysis Research Facility, AMMRF), as well as the PET facility at Lucas Heights, which will be able to use the developed technology to install similar remote CT services and become service providers.
3. Multiple research groups at Australian Universities and CSIRO which use CT as a tool in their research. Most of these groups are primarily interested in post-processing and analysis of CT reconstructed data with the purpose of extracting domain-specific information. An instructive example is presented by the work being carried out at CSIRO Petroleum Resources division in collaboration with several other CSIRO divisions, University of Western Australia, iVEC and others, aiming at extraction of elastic moduli, electrical resistivity and hydraulic permeability of geomaterials based on 3D micro-CT images. Similarly, the Department of Applied Mathematics and Research School of Physics and Engineering at ANU are world-renowned for their work on quantitative petrophysical analysis. They have developed powerful software libraries (MANGO, DRISHTI, etc.) which can be used for 3D image analysis, physics simulation and visualization. These tools require the results of CT reconstruction as input data, and as such they stand to benefit directly from the existence of remote CT reconstruction services that will provide CT data in standardized formats and with added flexibility with respect to the CT reconstruction parameters that can be readily fine-tuned for the needs of a particular form of subsequent analysis.

In the course of this project, we will develop generic libraries and toolkits for "web-enabling" and "cluster-enabling" of Windows desktop applications. These libraries can be used outside the CT context to provide the following functionality to a broad class of scientific desktop applications.

1. Distribution, automatic updating and controlled access to Windows applications, including a server-based user account management system.
2. Enabling a "manager" application running on user's computer to easily launch or queue a job on a remote compute cluster, passing any necessary parameters and receiving status information via HTTPS, and a "worker" application performing computation on the cluster to easily receive parameters and communicate status information to the "manager".

These libraries and tools can be reused to build other remote data and image processing services using uniform software architectural solutions. Prototypes of these libraries and tools have been developed by CSIRO in the last few years using standard ISO C++ and Microsoft libraries available in Visual Studio. Additional information about these systems (as well as on-line access to prototype services) can be found at the TS-Imaging website (http://www.ts-imaging.net/), and in particular at http://www.ts-imaging.net/WebHelp/X-TRACT/X-TRACT.htm.

Value adding

Identify the components of the project that could be based around generic technologies or be implemented through shared services for which the project would provide an exemplar use case or requirement set.

1. Data management component of the services will be based on the software and services being developed and deployed by ARCS, VeRSI, and the DataMINX project. Our prototype software is already capable working with standard simple image formats, such as TIFF, JPG, BMP and others, which can be presently used for importing and exporting CT image data in/from the system. If in the course of the proposed project it will be agreed (based on the consultations with major stakeholders and prospective users) that more advanced data formats should be used in order to alleviate long-term storage, query and retrieval of large quantities of the raw and reconstructed CT data, we will work with the DataMINX project to adopt standard data formats and management strategies developed there.
2. Data visualization component of the services will enable support of different 3D visualization software, possibly including DRISHTI (a free software developed by ANU, http://sf.anu.edu.au/Vizlab/drishti/), chiefly for tomographic data, running on specialized cluster nodes equipped with high-end GPU cards, in combination with VBL collaboration solution based on remote screen access.
3. Remote collaboration services will be overlayed (developed by VeRSI and the DataMINX project).
4. Data transfer solutions will be based on services developed by the DataMINX project.
5. AAA services to be integrated with existing solutions within the Australian Synchrotron with AAF compatibilities in mind for later integration.
6. The use of high speed network infrastructure leveraging the existing AARNet/AREN A3 backbone (may provide a user case for UCLP or GLIF based optical services)
7. Integration into 3D VIS environments one of which VeRSI is currently constructing. This will allow efficient and secure transmission of generated desktop images with 3D visualization over the Internet to client computers.

Standardisation

Describe the global technology development or standardisation work that will be adopted, adapted or extended within the project and any risk reduction available by collaboration with similar activities occurring else where in the world.

The participants of the present proposal are well aware of the state-of-the-art technologies for CT reconstruction, including those used at leading synchrotron facilities around the world.

We are presently working in collaboration with CT researchers from the European Synchrotron Radiation Facility, ESRF (Grenoble, France), Advanced Photon Source (Chicago, USA) and SPring-8 synchrotron (Harima, Japan).

The CT reconstruction software developed at CSIRO is being used at ESRF, Singapore Synchrotron Light Source, Lawrence Berkeley Laboratory in the US and elsewhere. We are also benchmarking our software against leading commercial CT software products such as Octopus (http://www.xraylab.com/) and Cobra (http://www.exxim-cc.com/products_cobra.htm).

In the area of remote CT systems, we have had very useful discussions with Francesco De Carlo and Brian Tieman from Advanced Photon Source (APS) in Chicago, who have developed and are running (at beamline 2BM at APS) what is, arguably, the most advanced remote CT system in the world. They are open for future collaboration, but still need to sort out some licensing issues in order to be able to share some of the software components. The available documentation also appears to be somewhat patchy at the present stage. We will use the experience accumulated at APS for building our version of remote CT services. It should be mentioned, however, that the APS implementation is not Windows-based, and is more of a collection of interoperable tools, rather than an integrated production system as proposed in our case. Both approaches have their merits and shortcomings, but the possibility for broad cross-fertilization between the two systems is partially limited by this fundamental difference in the underlying architecture.

Client/server communication for the remote computation framework is using standard SOAP web services. It is planned to implemented remote visualization using a standard VNC approach.

Project Scoping

Key Participants

Name any PfC components, any NCRIS capabilities, or any other institutions or groups that will need to be involved in the project planning and execution.

1. CSIRO Computational Imaging and Visualization project (Dr. Tim Gureyev, project leader)
2. CSIRO Computational and Simulation Sciences platform (Dr. John Taylor, platform leader)
3. VeRSI (Mr. Chris Myers, VBL Development Engineer)
4. Australian Synchrotron (Dr. Daniel Hausermann and Dr. Anton Maksimenko, Imaging and Medical Beamline, AS)
5. Dr Peter Turner, Project Manager, DataMINX project

Proposed members of the Project Committee:

1. Prof. Steve Wilkins (CSIRO), Leader of X-ray Science and Instrumentation group, (world-renowned specialist in advanced X-ray imaging methods, including Phase-Contrast CT). (Agreed)
2. Prof. Rob Lewis (Monash University), Director, Monash Centre for Synchrotron Science, (world-renowned expert in biomedical applications of synchrotron-based CT). (Agreed)
3. Assoc. Prof. Tim Senden (ANU), head of the X-ray CT group (world-renowned specialist in micro X-ray CT and its applications). (TBA)
4. Mr. Richard Farnsworth, Australian Synchrotron, Head of Controls and IT. (Agreed)
5. Dr. Ann Borda, Executive Director of VeRSI, or delegate. (Agreed)
6. Bill Appelbe, CEO and Chief Scientist of VPAC, or delegate (Agreed).
7. Executive Director of ARCS, Tony Williams, or delegate, currently Paul Coddington
8. Executive Director of ANDS, Ross Wilkinson, or delegate, currently Andrew Treloar

Project Scale

Identify the overall scale expected in the project, eg. 1 to 3 years, total effort in any year, and nominate any parties that have indicated a willingness to participate through providing resources. (funded or in-kind, people or facilities).

The proposed project term is 2 years, running at between 4.5 and 5 FTEs per year depending on the final budget and the stage of the project.

CSIRO will provide approximately 1.5 FTEs per year in-kind, as well as background IP and existing software modules for CT image reconstruction, software development infrastructure, as well as overall project management.

VeRSI will provide an equivalent of 1 FTE per year in-kind in the form of access to existing VBL and RLI solutions, expertise and skills in remote data access and management, remote collaboration services, AAA granular role based services, integration into existing infrastructure and will liaise with the DataMINX project.

The AS will provide this project with both conventional and computing resources as required for the development and testing of remote CT services, as well as continuing hosting of the service after the completion of the present project.

We are seeking the matching $300,000 per year funding from NeAT that will bring the total resources to the required equivalent of 6 FTEs per year.

If funding from NeAT is $250K per annum, it is anticipated that this would provide additional 1 FTE at VeRSI and 1 FTE at CSIRO. Funding of $300K per annum would also support 0.5 FTE at ANU, which would enable the remote tomography service to be deployed at the ANU CT facility, with some in-kind from ANU in the form of implementation of the services at the facility and comparison with results from ANU reconstruction codes running on the NCI National Facility..

Major Steps

Identify the key steps that will be visible to users as the services develop. Note that some observable deliverable is needed every half year and projects may be reviewed based on the achievement of these steps.

1. In the first 6 months we will develop a detailed architectural design of the remote CT services and identify computing and networking hardware requirements for implementing the solution. We will also engage with the proposed user community and provide them with remote access to the prototype CT service running on a small compute cluster at CSIRO in order to evaluate the user requirements and obtain feedback on aspects of the proposed solution.
2. In the second 6 months we will develop an initial basic solution for remote CT services and deploy it on the MASSIVE cluster at the AS. Access to this service will be limited to a small group of developers and expert users as part of the continuous development and testing cycle. During this time we will also develop basic infrastructure for remote visualization components of the services.
3. By the end of the third 6 months we will complete the development and deployment of the main software components, including the fully featured CT reconstruction and visualization components. We will complete basic integration of the data-management components using standardised tools and services, compatible with those proposed in the DataMINX project as available. At the end of this term, we will open access to the services based at the AS to a wider group of beta-testers. In this period, work would begin on development and testing of services at the second site.
4. The fourth 6 month period will be devoted primarily to testing and refinement of the remote CT services running at the AS, including load testing with very large (at least 4K^3 pixels) CT datasets and further debugging the software. We envisage some minor features to be added and modified at this stage as a result of feedback from the beta-testers. At this stage we will also complete the integration with the long-term data management and storage infrastructure. Through this last 6 months we will also be deploying and testing production versions of the services at the second site.

Dependencies

Identify dependencies that exist to activities or developments external to the project.

1. Deployment of the services will depend on the successful purchasing and commissioning of the computer cluster at the AS (the cluster is already fully funded and is scheduled to be purchased in 2009).
2. Data management and remote visualization components of the services will partially depend on the technologies developed by VeRSI in the course of other projects, including those proposed for DataMINX project.