

# **CSAR Service**

## **Consolidated Management Report**

### **2nd Quarter 2005**

#### **Management Summary**

This is the consolidated Management report for the second quarter 2005 of the CSAR HPC facility for UK Academia and Industry, which enables World-Class research and development.

The number of users stands at a total of 447 to date.

The workload on both the Origin 3000 Green and the Altix 3700 Newton has been fairly evenly spread across the mid- to high-end ranges of PEs during the second quarter of this year.

CSAR has been granted an 18-month extension of service contract until June 30<sup>th</sup> 2006. With this extension CfS implemented a further technology refresh introducing a 256 processor Itanium-2 (Madison) based SGI Altix, now enhanced to 512 processors with 1 Terabyte of memory.

CfS remains active in the UK Grid Forum.

## Introduction

This Management Report includes a section for each of the main service functions:

1. Service Quality
2. HPC Services
3. Science Applications Support Services
4. Training & Education Services
5. User Registration & New User Services
6. Value-Added Services

Each section includes a status report for the period, including notable achievements and problems, also noteworthy items for the next period.

## **1 Service Quality**

This section covers overall Customer Performance Assessment Ratings (CPARS), HPC System availability and usage, Service Quality Tokens and other information concerning issues, progress and plans for the CSAR Service.

### **1.1 CPARS**

Table 1 gives the measure by which the quality of the CSAR Service is judged. It identifies the metrics and performance targets, with colour coding so that different levels of achievement against targets can be readily identified. Unsatisfactory actual performance will trigger corrective action.

#### **CSAR Service - Service Quality Report - Performance Targets**

Service Quality Measure	Performance Targets					
	White	Blue	Green	Yellow	Orange	Red
<b>HPC Services Availability</b>						
Availability in Core Time (% of time)	> 99.9%	> 99.5%	> 99.2%	> 98.5%	> 95%	95% or less
Availability out of Core Time (% of time)	> 99.8%	> 99.5%	> 99.2%	> 98.5%	> 95%	95% or less
Number of Failures in month	0	1	2 to 3	4	5	> 5
Mean Time between failures in 52 week rolling period (hours)	>750	>500	>300	>200	>150	otherwise
<b>Fujitsu Service Availability</b>						
Availability in Core Time (% of time)	> 99.9%	> 99.5%	> 99.2%	> 98.5%	> 95%	95% or less
Availability out of Core Time (% of time)	> 99.8%	> 99.5%	> 99.2%	> 98.5%	> 95%	95% or less
<b>Help Desk</b>						
Non In-depth Queries - Max Time to resolve 50% of all queries	< 1/4	< 1/2	< 1	< 2	< 4	4 or more
Non In-depth Queries - Max Time to resolve 95% of all queries	< 1/2	< 1	< 2	< 3	< 5	5 or more
Administrative Queries - Max Time to resolve 95% of all queries	< 1/2	< 1	< 2	< 3	< 5	5 or more
Help Desk Telephone - % of calls answered within 2 minutes	>98%	> 95%	> 90%	> 85%	> 80%	80% or less
<b>Others</b>						
Normal Media Exchange Requests - average response time	< 1/2	< 1	< 2	< 3	< 5	5 or more
New User Registration Time (working days)	< 1/2	< 1	< 2	< 3	< 4	otherwise
Management Report Delivery Times (working days)	< 1	< 5	< 10	< 12	< 15	otherwise
System Maintenance - no. of sessions taken per system in the month	0	1	2	3	4	otherwise

**Table 1**

Table 2 gives actual performance information for the period. Overall, the CPARS Performance Achievement for the 2ndquarter 2005 was satisfactory (see Table 3), i.e. Green measured against the CPARS performance targets.

**CSAR Service - Service Quality Report - Actual Performance Achievement**

Service Quality Measure	2004/5											
	July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	Mar	April	May	June
<b>HPC Services Availability</b>												
Availability in Core Time (% of time)	97.97%	100%	99.52%	97.08%	98.50%	99.50%	97.37%	97.85%	97.85%	96.00%	99.50%	97.13%
Availability out of Core Time (% of time)	100%	99.2%	99.80%	98.67%	98.78%	99.2%	99.73%	99.5%	99.80%	99.90%	99.54%	99.22%
Number of Failures in month	2	2	2	3	4	2	3	5	4	4	1	2
Mean Time between failures in 52 week rolling period (hours)	5402	2345	1514	944	677	584	487	389	331	285	274	257
<b>Help Desk</b>												
Non In-depth Queries - Max Time to resolve 50% of all queries	<0.25	<0.25	<0.25	<0.25	<0.25	<0.25	<0.25	<0.25	<0.25	<0.25	<0.25	<0.25
Non In-depth Queries - Max Time to resolve 95% of all queries	<2	<1	<2	<3	<1	<0.5	<0.5	<2	<1	<2	<5	<1
Administrative Queries - Max Time to resolve 95% of all queries	<0.5	<0.5	<1	<1	<0.5	<1	<0.5	<2	<1	<1	<0.5	<2
Help Desk Telephone - % of calls answered within 2 minutes	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%	100%
<b>Others</b>												
Normal Media Exchange Requests - average response time	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5	<0.5
New User Registration Time (working days)	0	0	0	0	0	0	0	0	0	0	0	0
Management Report Delivery Times (working days)	10	10	10	10	10	10	10	10	10	10	10	10
System Maintenance - no. of sessions taken per system in the month	2	2	2	2	2	2	2	2	2	2	2	2

**Table 2**

Notes:

- HPC Services Availability has been calculated using the following formula, based on the relative NPB performance of Fermat, Green and Newton at installation:  

$$[\text{Fermat availability} \times 40 / (40+233+343)] + [\text{Green availability} \times 233 / (40+233+343)] + [\text{Newton availability} \times 343 / (40+233+343)]$$
- Mean Time Between Failures for Service Credits is formally calculated from a rolling 12-month period.

Table 3 gives Service Credit values for each month to date. These are accounted on a quarterly basis, formally from the Go-Live Date. The values are calculated according to agreed Service Credit Ratings and Weightings.


**CSAR Service - Service Quality Report - Service Credits**

Service Quality Measure	2004/5											
	July	Aug	Sept	Oct	Nov	Dec	Jan	Feb	Mar	April	May	June
<b>HPC Services Availability</b>												
Availability in Core Time (% of time)	0.078	-0.058	-0.039	0.078	0.039	-0.039	0.078	0.078	0.078	0.078	0	0.078
Availability out of Core Time (% of time)	-0.047	0	-0.047	0	0	0	-0.039	0	-0.047	-0.047	-0.039	0
Number of Failures in month	0	0	0	0.008	0.008	0	0.008	0.0004	0.008	0.008	-0.008	0
Mean Time between failures in 52 week rolling period (hours)	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002
<b>Help Desk</b>												
Non In-depth Queries - Max Time to resolve 50% of all queries	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019
Non In-depth Queries - Max Time to resolve 95% of all queries	0	-0.016	0	0.016	-0.016	-0.019	-0.019	0	-0.016	0	0.046	-0.016
Administrative Queries - Max Time to resolve 95% of all queries	-0.019	-0.019	-0.016	-0.016	-0.019	-0.016	-0.019	0	-0.016	-0.016	-0.019	0
Help Desk Telephone - % of calls answered within 2 minutes	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004	-0.004
<b>Others</b>												
Normal Media Exchange Requests - average response time	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002	-0.002
New User Registration Time (working days)	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019	-0.019
Management Report Delivery Times (working days)	0	0	0	0	0	0	0	0	0	0	0	0
System Maintenance - no. of sessions taken per system in the month	0	0	0	0	0	0	0	0	0	0	0	0
Monthly Total & overall Service Quality Rating for each period:	-0.02	-0.07	-0.07	0.02	-0.02	-0.06	-0.02	0.02	-0.02	-0.01	-0.03	0.01
Quarterly Service Credits:	-0.04		-0.12		-0.09		-0.01					

**Table 3**


### 1.2 No. of Registered Users

The current position at the end of the quarter is that there are 447 registered users of the CSAR Service.



### 1.3 Service Quality Tokens

The graph below illustrates the monthly usage trend of Service Quality Tokens:



Over the course of the quarter the position is that as a management tool the Service Quality Tokens have been available to enable the users to provide qualitative feedback about all aspects of the service. This feedback is used as a mechanism to initiate change in the service where appropriate.

At the end of the quarter no gold stars or black marks had been allocated to the service.

## 2 HPC Services Usage

Usage information is given in tabular form, and in graphical format. The system usage information covers:


- CPU usage
- User Disk allocation
- HSM/tape usage

This is illustrated in a number of graphs including:


- a) SMP (Altix/Origin) Usage by month, showing usage each month of CPU (GFLOP-Years as per NPB), split by Research Council and by system. Overlaid horizontal lines show the overall Capacities.
- b) SMP (Origin) Usage by month, showing usage each month in CPU Hours, split by Research Council and giving the equivalent GFLOP-Years as per NPB. The Baseline and overall Capacity are shown by overlaid horizontal lines.
- c) Medium Performance Disk (Origin) allocated for User Data by month, showing the allocated space each month in GBytes, split by Research Council. The Baseline Capacity (1.5 Terabytes) is shown by an overlaid horizontal line.
- d) HSM/Tape Usage by month, showing the volumes held each in GBytes, split by Research Council. The Baseline Capacity (16 Terabytes) available will be shown by an overlaid horizontal line.


## 2.1 Service Usage Charts

The graphs below show recent monthly CPU, disk and HSM allocations and usage.




Usage of the SGI Altix 3700 system Newton is shown by Research Council in the above graph.







Usage of the two batch SGI Origin systems, Fermat and Green, is shown by Research Council during the last 12 months of service in the preceding two charts.




The above chart, and the one below, shows the wait time trend in hours on the Altix 3700 Newton.



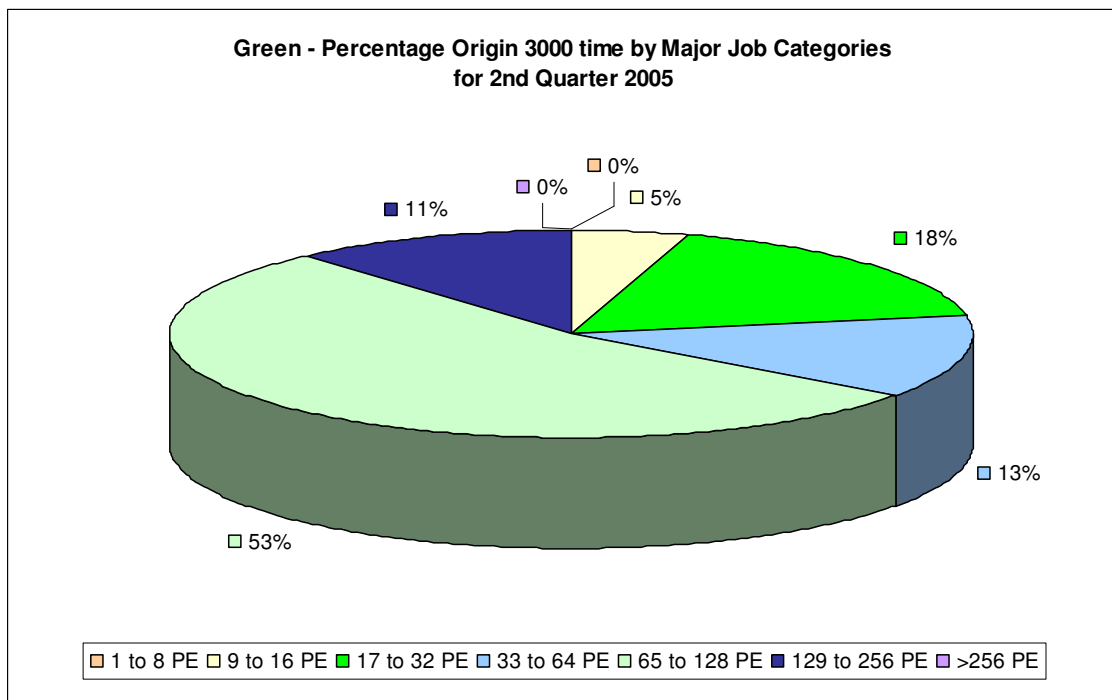
The following two charts show average wait times in hours for the quarter on the Origin 3000 Green.









The next series of four charts show the percentage CPU time utilisation by the major job categories on the Green, Fermat and Newton systems for the 2nd quarter 2005.




During this quarter there has been a good spread of work across the PE ranges on Green.




The workload on Fermat during this quarter was concentrated primarily in the mid-range PEs.




There was a good spread of work across Newton during the second quarter.



Usage on Green tended more to be in the mid-range PEs during this quarter.



The workload on Fermat for this quarter was spread fairly evenly across the mid-range PEs.




During the second quarter usage on Newton was concentrated mainly in the mid- to high-range PEs.

## 2.2 System Usage Graphs


In all the Usage Charts, the baseline varies dependant upon the number of days in each month, within a 365-day year.

### 2.2.1 Baseline System


The graph below shows the Gflop Years utilisation on the CSAR systems by Research Council for the last 12 months.





The graph below shows the same service utilisation by major system.




The next chart shows the historic interactive usage of the Origin 300 system Wren. Eight of the higher speed 500Mhz CPUs in Wren deliver the baseline capacity equivalent to that which was previously available on the Origin 3000 system Fermat for interactive usage.



The following two charts detail the historic usage of the Origin 3000 system (Green) and the Altix 3700 system (Newton).




The next series of graphs illustrates the usage of the Medium Performance disk and HSM/tmp resources of the system.



The above graph illustrates the historic allocation of the Medium Performance Disk on Fermat and the SAN.

The next graph shows the historic HSM usage by Research Council funded projects, which has exceeded the overall Baseline of 16 Terabytes, and now totals about 48 Terabytes.



### 2.2.2 Guest System Usage

There is currently no Guest System usage.

### 2.3 Capability Incentives

Capability incentives were historically given on the T3E system Turing for jobs of 512 PEs and above. In July 2003 it was announced that discounts for capability jobs available on all CSAR systems had been approved to include the SGI Origin 3000 system Green and the SGI Altix 3700 system Newton.

These capability incentives were agreed with the Research Councils to encourage capability usage of the national supercomputers for greater scientific achievement, and offer the following discounts:

System	No of Processors	Discount
newton	192+ CPUs	15% discount
newton	128+ CPUs	10% discount
green	384+ CPUs	15% discount
green	256+ CPUs	10% discount

Discounts are given in the form of refunded Service Tokens.

Changes in usage patterns will be monitored and, subject to review, CfS reserve the right to change the incentives at any future date.

The following table displays the capability incentive discounts granted during the second quarter of this year.

Service Tokens Refunded: Quarter 2 2005 Usage							
System		Consortia					Total
		cse075	cse086	cse121	csn006	csn015	
Green	256+ PEs					89.99	89.99
Green	384+ PEs						0
Newton	128+ PEs	239.6		240.34	51.76		531.7
Newton	192+ PEs	574.85	143.93				718.78
<b>Total Tokens</b>							<b>1340.47</b>

This is within the CfS Management Board's forecast.



## **2.4 Service Status, Issues and Plans**

### **Status**

The service has been reasonably utilised throughout the second quarter of 2005, with usage exceeding baseline.

There was a relatively balanced spread of work across all major systems throughout Quarter 2.

In the final week of April, the operating systems on both the Origin and the Altix systems were upgraded, along with the associated CXFS SAN software. Problems experienced during previous months with the CXFS software had been targeted to be fixed by the vendor in this new release, and it was hoped that all previous issues have now been addressed with this upgrade.

Following these upgrades, the stability of all systems, most notably the Altix systems, did greatly improve during the second quarter, with few outages encountered.

### **Issues**

There were some periods of system downtime during the early part of the quarter, primarily affecting the Altix system Newton, which were subsequently addressed with the upgrades as mentioned above.

### **Plans**

It is planned to combine the two 256 PE Newton nodes into one 512 processor physical node, expected to be carried during the early part of the third quarter this year. Additionally at the same time, it is planned to convert the newly-joined system into a 512 PE Single System Image, which will permit jobs to be run of a greater size than is currently possible.

### **3 Project Management, Documentation and User Feedback**

This section covers aspects relating to the registration of projects and users, the management of projects and resources, topics associated with documentation and user feedback.

#### **3.1 Project Applications**

21 applications for new CSAR projects were received, requesting a total of 727,906 service tokens.

#### **3.2 New Projects**

6 new CSAR projects were started with 71,268 service tokens being awarded in total.

#### **3.3 Finished Projects**

2 projects finished

#### **3.4 Queries**

A total of 165 CSAR queries were dealt with:

- o 115 non-in-depth
- o 40 user registration and admin
- o 10 in-depth:

#### **3.5 Service Quality Tokens**

No gold stars or black marks were awarded during this period.

#### **3.6 CSAR Focus**

The 13<sup>th</sup> edition of CSAR Focus was published and is also available online at [www.csar.cfs.ac.uk/about/csarfocus/focus13/focus13.pdf](http://www.csar.cfs.ac.uk/about/csarfocus/focus13/focus13.pdf).

## **4 Scientific Application Support Services**

### **4.1 Training and Education**

The following courses were delivered:

- Itanium2 and Parallel Scaling for the SGI Altix 3700, University of Manchester
- Performance and Development Tools, University of Manchester
- Jon Gibson delivered the course “Introduction to MPI” at the Mathematical Institute, St Andrews University
- Kevin Roy delivered an introductory course on using CSAR and HPC at QMUL

### **4.2 Consortia Support/Software**

The latest version of the SGI Operating System, Propack 3.0, was installed on Newton. The default version of the compilers was also changed to version 8.1 release 24. Other work carried out includes:

#### **4.2.1 h2mol**

Craig Lucas and Kevin Roy visited Queen’s University Belfast to discuss CSAR running the h2mol code on their behalf. This work has now started.

#### **4.2.2 Data Visualisation**

CSAR staff have worked on implementing a library of routines to conveniently write out data to be input into AVS. This software works in parallel taking advantage of MPI-IO and Parallel HDF.

#### **4.2.3 Amazon-SC**

Work has started on optimising and parallelising the Amazon-SC code (for the numerical simulation of wave overtopping at seawalls) on behalf of a project based at Manchester Metropolitan University.

#### **4.2.4 VASP**

A new SGI-optimised version of VASP has been tested and installed on Newton. As well as the full version of VASP 4.6.26, a gamma-point only version has been installed. The new version shows a significant improvement in runtimes and scaling over the previously installed VASP 4.6.7.

#### **4.2.5 Netsolve**

The NetSolve group have been assisted with the developing of the GridSolve software so that it can potentially be used on CSAR machines. GSI authentication has been added. A new version is now available to be tested on CSAR machines.

#### **4.2.6 Tomcat**

Work has continued on the development of a dynamic load balancing scheme for Tomcat.

#### **4.2.7 DL\_POLY**

The latest versions of DL\_POLY2 and DL\_POLY3 have been installed.

#### **4.2.8 Accounting**

A monthly accounting script which enables PIs to get an instant view of token usage for a given month at the user level, and a tape resource script “holdsize” (and variant “tapesize”) which gives tape usage information for a given user have been created.

**4.2.9 ns2 (Network Simulator 2)**

NS2, the Network Simulator code version 2, which is used to model and simulate the behaviour of the computer networks, was installed on the Origin systems and on Newton, on behalf of a “New User” project under Prof. Pitt’s at Queen Mary University.

**4.2.10 NetCDF Operators (nco)**

The latest version of this package, which is used for the manipulation and analysis of data stored in netCDF format, was installed on the Origin systems. This is used mainly by the NCAS Consortium.

## 5 Collaboration and Conferences

### 5.1 MRCCS Projects

#### 5.1.1 Virtual Prototyping

The virtual prototyping interactive FEA library has been integrated with RealityGrid to create "RapidFire". The work is currently being extended to demonstrate the potential for dinosaur footprint simulation.

An abstract for the TCN CAE conference in Italy has been submitted: Margetts, L., Ford, R., Porter, A. and Haines, R. 'RapidFire: A novel grid enabled application for interactive finite element analysis'.

An extended abstract for Margetts, L., Smith, I.M., Leng, J., 'Simulating Dinosaur Trackway Formation', has been submitted for COMPLAS, Barcelona, September 2005.

#### 5.1.2 MRCCS Seminars

The Access Grid based seminar series on HPC and visualisation related topics continued on Fridays, including:

- Henggui Zhang, University of Manchester – Virtual Tissue: Engineering the Mammalian Heart from Cell to Bedside
- Karen Grainger, the Stereoscopic Society – Curating and the Collection: Future Plans for a Digital Stereo Image Archive and Touring Exhibitions in 3D
- Simon McIntosh-Smith, Clearspeed – The Next Wave: co-processors for application acceleration

### 5.2 Events

Events attended include:

- MOLPRO workshop, Imperial College London, 5<sup>th</sup> – 6<sup>th</sup> April
- UK-KDD05 Workshop, Liverpool, 6<sup>th</sup> April
- Lee Margetts gave an invited seminar "Terascale Simulation using Byte-size Pieces" to the MHD Research Group, University of Exeter, April 25<sup>th</sup>
- DL\_POLY training day, Daresbury Laboratory, 27<sup>th</sup> April
- NAFEMS World Congress, Malta, May 2005. Lee Margetts presented the paper: Margetts, L., Smethurst, C., Ford, R., 'Interactive Finite Element Analysis'
- Cray User Group, 16<sup>th</sup> – 19<sup>th</sup> May
- NEC User Group XVII, Exeter, 25<sup>th</sup> and 26<sup>th</sup> May
- SGI User Group, 13<sup>th</sup> – 16<sup>th</sup> June. Kevin Roy presented two papers at this conference: *Effective Bandwidth Utilisation of multi-user systems* and *Dynamic load balancing using single sided communication*

### 5.3 Publications

Cooper, P.D., Kjaergaard, H.G., Langford, V.S., McKinley, A.J., Quickenden, T.I., Robinson, T.W., and Schofield, D.P. Infrared identification of matrix isolated H<sub>2</sub>O.O<sub>2</sub>, J. Phys. Chem. A 109 (19) 4274-4279, 2005

Mills, M.J., Toon, O.B., Vaida, V., Hintze, P.E., Kjaergaard, H.G., Schofield, D.P., and Robinson, T.W. Photolysis of sulfuric acid vapor by visible light as a source of the polar stratospheric CN layer, J. Geophys. Res.(Atmos), 110 (D8): art. no. D08201, 2005-09-06

Margetts, L. 'Engineering Simulation – Visions of the Future', UniIT newsletter, University of Manchester

Smith, I.M. and Margetts, L. 'The convergence variability of parallel iterative solvers' which has been accepted for publication in Engineering Computations

## **6 Added Value Services**

### **6.1 ISC 2005**

Three members of Manchester Computing attended and exhibited at the International Supercomputer Conference (ISC), which took place in June (21<sup>st</sup>-24<sup>th</sup>), and this year the event celebrated its 20<sup>th</sup> anniversary.

At the exhibition stand, demonstrations were given on the Parallel System Toolkit (PST), for parallel modules in AVS/Express and RapidFire, an FEA application employing real-time computation on parallel systems in a virtual environment.

### **6.2 SC2005**

Preparations are underway for Manchester Computing's exhibition at this year's Supercomputing Conference to be held in Seattle in November.

### **6.3 Visualisation**

The SAGE project at MC has created a portable Virtual Environment system, using a standard PC and commercial off-the-shelf products. The project has produced a description of how to build the system from components, available from <http://kato.mvc.mcc.ac.uk/sve-wiki/SAGE> so that research departments, from engineering to geology, from design to performance art, can build their own for between 7 to 10 thousand pounds. A fully working version of the system is currently being demonstrated at various events across the country. An article in the last CSAR Focus newsletter gave the benefits of stereoscopic projection, a technology used in these systems to allow full 3D vision of visualisation scenes. We welcome interest from the CSAR community.

### **6.4 National Grid Service**

Work has been completed assisting a research group on the use of the NGS service machines. This included getting E-Science certificates, setting up Globus to get a proxy, compiling and running codes. This research group has also applied for CSAR resources.