HPCx Quarterly Report
July – September 2007

1 Introduction

This report covers the period from 1 July 2007 at 0800 to 1 October 2007 at 0800.

The next section summarises the main points of the service for this quarter. Section 3 gives details of the usage of the service, including failures, serviceability, CPU usage, helpdesk statistics and service quality tokens. A summary table of the key performance metrics is given in the final section. The Appendices define the incident severity levels and list the current HPCx projects.

2 Executive Summary

- This was another exceptional quarter for reliability. We have now not had a single failure for over six months.
- Utilisation has dropped somewhat from the high levels seen last quarter, presumably due to the effects of the summer break. Utilisation for the development service continues to be very high at over 90%.
- We remain on target both for technical reports and training. A visualisation workshop was run this quarter, and we are currently taking registrations for both the second technical workshop and the Annual Seminar.
- Work by HPCx technical staff was presented at a number of international meetings including two talks at ScicomP13 in Garching and one at ParCo2007 in Juelich. One of these ScicomP talks arose directly from work done for last quarter's technical report HPCxTR0703.
- HPCx staff participated in and gave talks at consortia-related workshops organised by UKTC and by CCP2.
- The terascaling work on the OCCAM code was successfully completed and resulted in the award of a gold capability incentive.
- A new release of DL_POLY_3, version 08, was released after substantial user testing. It contains many important changes from the previous version.

- Comparative performance studies on IBM POWER and Cray XT architectures are continuing with the aim of publishing a paper. Initial IO results for these platforms are given in the technical report HPCxTR0707.
- The Software Engineering team's activities in HPC tools and languages include Dr Mark Bull's chairmanship of the OpenMP Architecture Review Board's Language Committee. This committee has just released the draft 3.0 standard for public comment, the first language update in over two years.
- The Grid middleware support on HPCx has been extended via installation and investigation of the most recent releases of Globus.
- An industrial discrete-element modelling application from DEM Solutions Ltd. has been ported to HPCx. We now plan to perform scaling studies to see if there is a business case for them to purchase HPCx cycles to perform their most challenging simulations.

3 Usage Statistics

3.1 Availability

3.1.1 Failures

The monthly numbers of incidents and failures (SEV 1 incidents) are shown in the table below:

	July	*August*	*September*
Incidents	6	0	4
Failures	0	0	0

Thus, there were no failures this quarter.

3.1.2 Performance Statistics

This section uses the definitions agreed in Schedule 7, ie,

- MTBF = (24 x 30.5)/(number of failures in month)
- Serviceability (%) = 100 x (WCT – SDT – UDT) / (WCT– SDT)

Attribution	*Metric*	*July*	*August*	*September*	*Quarterly*
IBM	Failures	0	0	0	0
	MTBF	∞	∞	∞	∞
	Serviceability	100.0%	100.0%	100.0%	100.0%
Site	Failures	0	0	0	0
	MTBF	∞	∞	∞	∞
	Serviceability	100.0%	100.0%	100.0%	100.0%
External	Failures	0	0	0	0
	MTBF	∞	∞	∞	∞
	Serviceability	100.0%	100.0%	100.0%	100.0%
Total	Failures	0	0	0	0
	MTBF	∞	∞	∞	∞
	Serviceability	100.0%	100.0%	100.0%	100.0%

3.2 Utilisation

The graphs below show the overall utilisation of the two services, and the proportion of the main service utilisation which was classed as capability work – that is, jobs which used more that 256 processors.

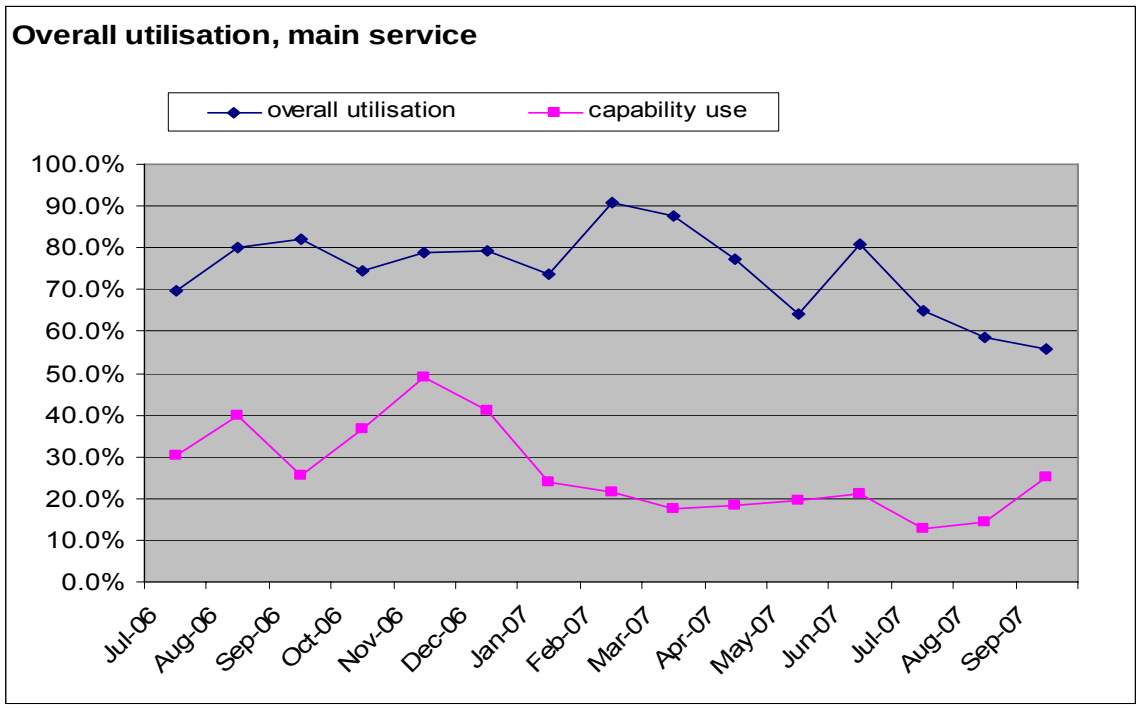

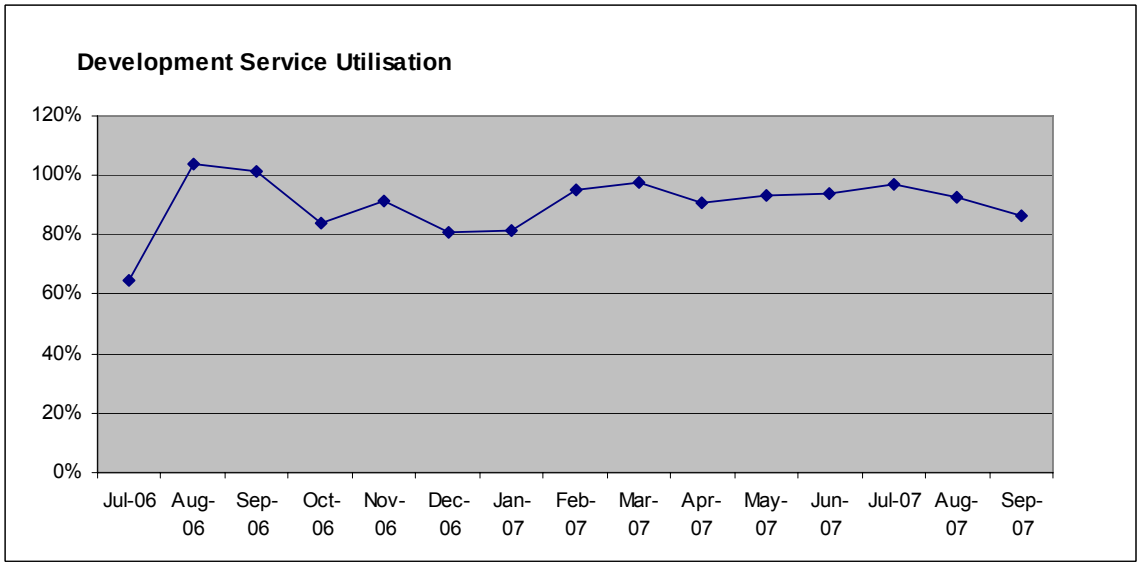

Utilisation figures greater than 100% for the development service correspond to a period in August and September 2006 when the number of processors in the service was temporarily increased.

3.3 Capacity Planning

Predicted Utilisation

The graph below shows the utilisation since the start of the project and the projected utilisation (on the main service) until January 2008. The scale on the *y*-axis is AUs per hour, where at peak Phase 3 can deliver 12034 AUs per hour (the upper red line in the graph). The lower line (in blue) corresponds to the more practicable 80% level.

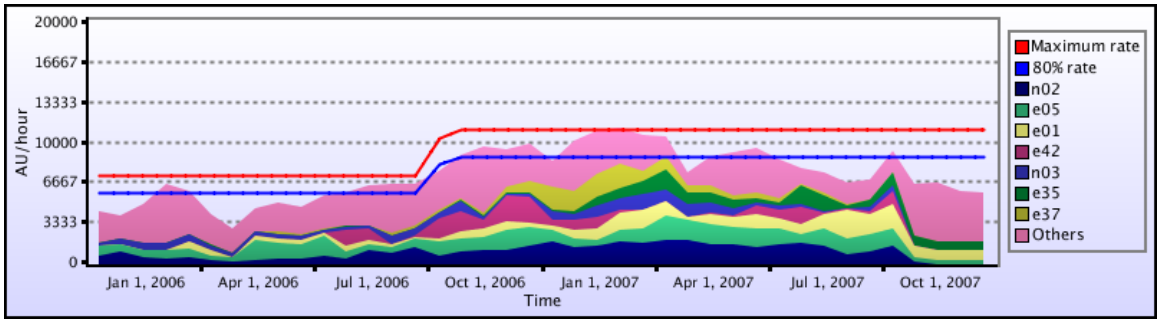

The graph assumes:

- that each project will use its remaining allocation pro rata with its usage profile as known to the database, which is often simply that on the original application form;
- that no more projects are given access to the service.

The graph shows that, on the basis of the projects which are currently using the service, we can anticipate a little spare capacity later in 2007.

Numbers of Research Consortia

At the end of this quarter there were 42 research consortia on HPCx. In addition, there is one active externally funded project.

3.4 CPU Usage by Job Size

Main service

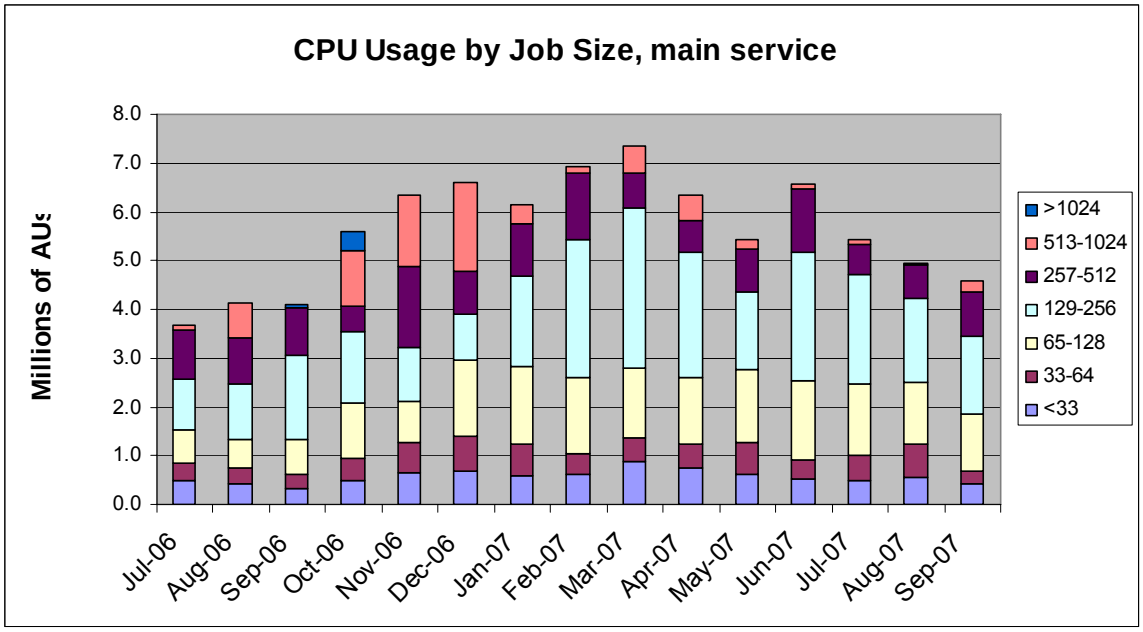

Development Service

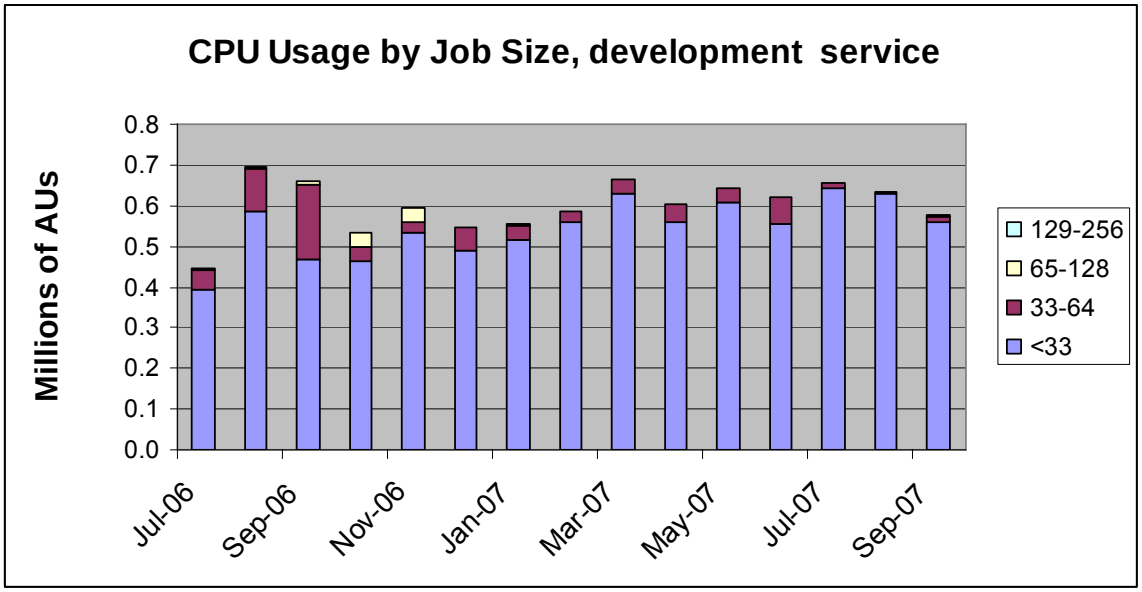

3.5 AU Usage by Consortium

Main Service

Consortium	*July*	*August*	*September*	*AUs charged*	*%age of charged AUs*
e01	582872	817761	1507157	2907790	19.6%
e03		5		5	0.0%
e05	815882	867486	1035496	2718864	18.3%
e06	9707	4906		14613	0.1%
e08	113236	122013	56485	291734	2.0%
e10		3347		3347	0.0%
e11	50744	39586		90330	0.6%
e17		1295	1584	2879	0.0%
e18	9822	11804	4035	25661	0.2%
e19			1442	1442	0.0%
e24	601030	4417	284	605731	4.1%
e25	1907			1907	0.0%
e26		6052		6052	0.0%
e31	52327			52327	0.4%
e33	267530	625043	322875	1215448	8.2%
e35	391371	1026513	334222	1752106	11.8%
e36		5614	17574	23188	0.2%
e37	189042	72373		261415	1.8%
e38	4723	3896		8619	0.1%
e39	58153	83076	95476	236705	1.6%
e41	650	31813	7456	39919	0.3%
e42	783952	214094	167639	1165685	7.9%
e44		4	1	5	0.0%
e45			74103	74103	0.5%
e46	3889	7643		11532	0.1%
e49	3649	20092	53680	77421	0.5%
e50	68683	14363	40218	123264	0.8%
e51	159			159	0.0%
e53	4910	31417	16134	52461	0.4%
e59	907		7381	8288	0.1%
e60		31171	75983	107154	0.7%
e61	191478	2857	18429	212764	1.4%

e62	35	10,146	5705	15886	0.1%
e63			17325	17325	0.1%
EPSRC Total	4206658	4058787	3860684	12126129	81.8%

n01	234948	1255	13536	249739	1.7%
n02	595269	538447	171308	1305024	8.8%
n03	227612	118391	159311	505314	3.4%
n04	14045	110541	251487	376073	2.5%
NERC Total	1071875	768634	595642	2436150	16.4%

p01	7699	117	11446	19262	0.1%
PPARC Total	7699	117	11446	19262	0.1%

c01	9112	2168	13326	24606	0.2%
CCLRC Total	9112	2168	13326	24606	0.2%

b08	62141	36404	5661	104206	0.7%
BBSRC Total	62141	36404	5661	104206	0.7%

x01	8342	98	4037	12477	0.1%
External Total	8342	98	4037	12477	0.1%

z001	11834	33667	44066	89567	0.6%
z004	9288	8529	1614	19431	0.1%
z06		19		19	0.0%
HPCx Total	21124	42229	45681	109034	0.7%

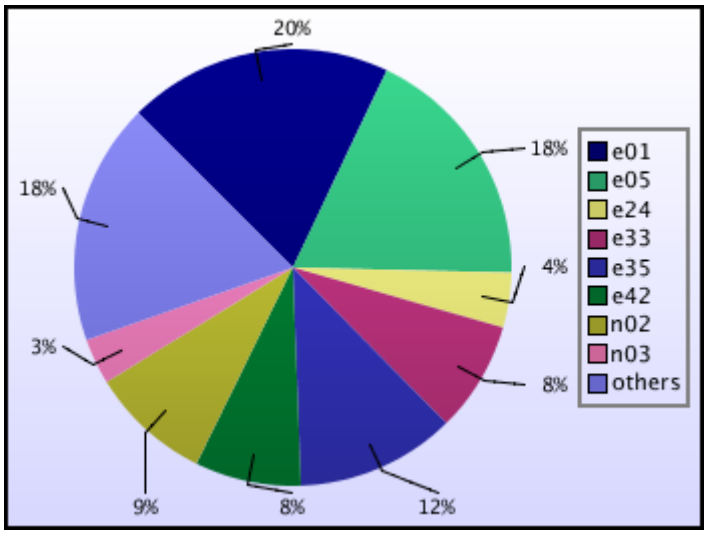

Development service

Consortium	*July*	*August*	*September*	*AUs charged*	*%age of charged AUs*
n01	16	0	0	16	0.0%
n02	645511	631063	560052	1836626	97.7%
n03	561	3859	8,204	12624	0.7%
n04	21436		9190	30626	1.6%
NERC total	667524	634922	577446	1879892	100.0%

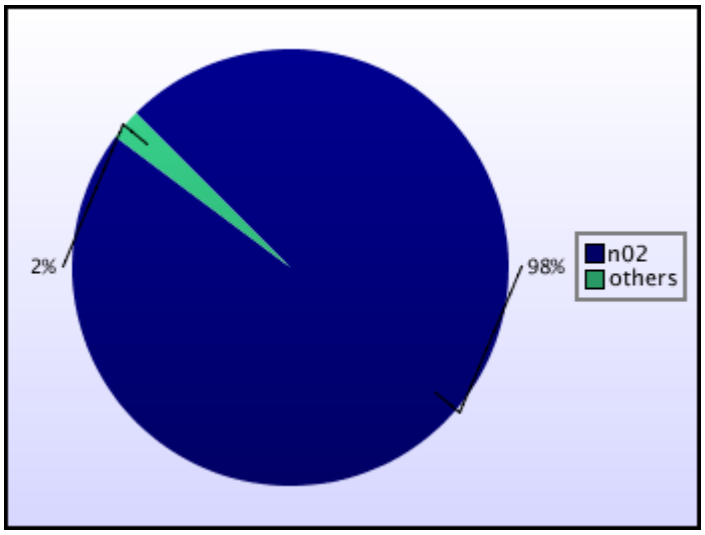

3.5.2 Discounts done

The following table shows the discounts that were awarded during the last quarter.

Consortium	*AU used*	*AU charged*	*Discount*
e01	2965440	2907789	57650
e05	2759509	2718864	40645
e36	27279	23187	4092
n03	532752	505314	27438
n04	397396	376074	21322
e01	2965440	2907789	57650

3.6 Helpdesk

3.6.1 Classifications

Category	*Number*	*% of all*
Administrative	79	41.6%
Technical	91	47.9%
In-depth	15	7.9%
Technical Assessment	4	2.1%
PMR	1	0.5%
Total	190	100.0%

3.6.2 Performance

All non-indepth queries	*Number*	*%*	*Target*
Finished within 24 Hours	118	69%	75%
Finished within 72 Hours	165	97%	97%
Finished after 72 Hours	25		

Administrative queries	*Number*	*%*	*Target*
Finished within 48 Hours	72	91%	97%
Finished after 48 Hours	7		

3.6.3 Experts Handling Queries

Expert	*Admin*	*Technical*	*In-Depth*	*PMR*	*Technical Assessment*
sysadm	28	32	4	1	0
DL	3	23	4	0	3
EPCC	46	36	7	0	1
Other	2	0	0	0	0

3.7 Service Quality Tokens

Sep 28, 2007 8:48:20 AM	Dr Daniel Mason	* * *

4 Support

4.1 Applications Support *(Dr David Henty)*

4.1.1 Documentation

The online copies of the IBM manuals have been updated to reflect the recent upgrades to version 10 of the compiler suite. The Totalview documentation also had to be modified after installation of the current product release.

4.1.2 Technical Reports

Two reports were planned for Q3 in the following areas:

a) Novel HPC Languages

b) IO Performance on HPCx

We have produced the following three reports this quarter:

- **HPCxTR0706** *Fluent on HPCx,* A.G. Sunderland.
- **HPCxTR0707** *Chapel, Fortress and X10: novel languages for HPC,* M. Weiland.
- **HPCxTR0708** *DL_POLY_3 I/O Analysis, Alternatives and Future Strategies*, I.J. Bush and I. Todorov.

Reports 07 and 08 directly address the two planned topics a) and b); report 06 resulted from applications support work done for an HPCx consortium.

4.1.3 Training

There were no training courses run during Q3: a planned run of an MPI course in Leeds was rescheduled due to issues with availability of a suitable venue. Statistics are summarised below alongside annual targets (where appropriate):

Metric	Total	Target
Course days	13	20
Different course titles	5	6
Different locations	2	4
Student-days for HPCx users	167	
Student-days for HPCx staff	27	
Student-days available for HPCx	311	

We have already arranged all but one of the required seven days training in Q4:

- one-day DL_POLY course at Daresbury in late November (alongside the Annual Seminar);
- two-day course on High-Performance Reconfigurable Computing (eg FPGAs) at EPCC in early December;
- three-day Message-Passing course in Leeds on 12-14 December.

4.1.4 Workshops and Conferences

The second of the two workshops for this year addresses the practical use of HPC tools for capability computing, and will be held at RAL on 5 November; we are already taking registrations for this workshop via the HPCx WWW pages. The main conference for the year, the Fifth Annual HPCx Seminar, will be held in Daresbury on 26 November alongside the 18th Machine Evaluation Workshop. Preparations for this event are well in hand with the majority of speakers already confirmed.

4.1.5 User Group

The first user group meeting for the year was arranged using Access Grid on Thursday 16th August. Unlike previous User Groups held over AG, this meeting was not a success with practically no external users in attendance. We will therefore need to consider whether we should continue to hold meetings in this manner, or whether they should all be face-to-face events. The second and final User Group for the year will take place immediately after the Annual Seminar in November, and we would expect many more users to attend this meeting.

4.1.6 Newsletter

Issue eight of Capability Computing is due to go to press in the second week of October, in time for distribution at Supercomputing 2007.

4.2 Outreach Activities *(Dr Richard Blake)*

Progress against key objectives:

4.2.1 Life Sciences

With the completion of the Life Sciences funding, there will be a limited level of resources available for Outreach. Major activities for 2007 will be:

- demonstration of the retina modelling code on much larger data sets.

Work is still ongoing in terms of analysing sequential bottlenecks in the inclusion of the image field.

Discussions are underway with BBSRC with a view to establishing two focussed workshops on Cell and Physiome Simulation. These are being progressed as part of the development of the Hartree Centre – consultation is underway with Peter Coveney and Charles Laughton as to participants and programme.

4.2.2 Public/Industrial Awareness

We will aim at improving public and industry awareness, in particular through engagement with Science Festivals and marketing activities:

- continued efforts to get funding for a longer term Public Understanding of Science programme around HPC;

Over the summer EPCC employed a student to look at developing simple graphical demonstrations to explain the use of HPC to the general public. This work was quite successful, resulting in a review of existing technologies for producing visualisations and a prototype visualisation of a parallel traffic model.

Unfortunately, our bid to EPSRC for a Partnerships for Public Engagement grant to fund a coordinated programme of HPC outreach activities was not funded.

- ongoing promotion of access to HPCx to industry as part of the general marketing activities at EPCC and CCLRC.

A commercial parallel code for Discrete Element Modelling has been ported to HPCx and verified on small test cases. Next quarter we plan to investigate scaling of bigger datasets to many processors, with the future possibility of selling cycles to DEM Ltd to perform simulations too large for their local cluster.

4.2.3 Exploiting the CCLRC-funded visualisation facility

We plan to enable HPCx users to take advantage of this visualisation resource, which is physically collocated with the HPCx system. To this end we will:

- develop two new collaborations with high-end computing Consortia;

Visualisation projects are still being explored with the Proudman Oceanographic Laboratory and the plasma fusion simulation group at Culham. These discussions are focussing on the software requirements and potential scenarios for multi-dataset visualisation.

- hold a Workshop on High-end Visualisation in Engineering in collaboration with Viznet.

We held a Visualization workshop at Daresbury 24th September 2007 under the HPCx and vizNET banners. The title of the workshop was "Advanced Visualization for High-End Computing". The web page for the workshop is at www.cse.scitech.ac.uk/events/Visualisation2007/*. Some 26 people registered.*

- support for demonstrations of HPCx integration into DEISA;

10 Gb/s link procured and operational in unidirectional mode. Bi-directional issues being investigated by Net-North West.

- support for joint EPSRC-NSF collaborations with Teragrid

GENIUS grant proposal to simulate cranial blood flow in a Grid environment funded and currently exploring implementation of advanced reservation and MPIg implementation.

4.3 Terascaling Applications *(Dr Mike Ashworth)*

4.3.1 Key Objectives

Progress against key objectives:

Publications
Three technical reports have been produced this quarter: a summary of benchmark results from the Fluent engineering package, a comparative description of the new parallel languages to emerge from the HPCS Programme and a report on evaluation of I/O performance of the DL_POLY code.

Performance of codes on future systems
J. Hein, A. Gray, F. Reid, A. Simpson, M. Bull, M. Weiland and S. Booth have continued to port a number of codes to HECToR, many of which are currently being used on HPCx. They have also been measuring the performance of these codes on HECToR, in order to produce results for an upcoming performance report comparing performance to HPCx.

Interactions with Consortia:
A. Gray and K. D'Mellow visited the e40 (Computational Quantum Many-Body Theory) consortium researchers at University of York on 27th September 2007. This visit was successful in strengthening our collaboration and furthering the parallelisation and optimisation of the GWST code.

M. Plummer and A.G. Sunderland attended and presented at the R-Matrix Computations Workshop, UCL, 14th-15th September 2007.

M. Ashworth, D.R. Emerson and A.G. Sunderland attended the UK Turbulence Consortium Workshop, University of Southampton, UK, 26th-27th September 2007. M. Ashworth gave a presentation on "Petascale Computing".

4.3.2 Computational Materials Codes

CASTEP: CASTEP has been debugged on an IBM BlueGene/L system by M. Plummer and K. Refson: certain system-call routines which caused crashes have been rewritten to be correct 'iso_c_binding' F2K3 Fortran (allowed by the F95 compiler). These corrections and other uses of 'iso_c_binding' are being fed back for testing on HPCx CASTEP and the Cray version. Certain limitations of the BlueGene compiler have been found and reported to IBM.

A fault in the IBM XLFortran compiler which can cause slow-down of certain CASTEP jobs and in some cases unnecessary memory overload has been reported to IBM as an HPCx PMR. This fault occurs when pointer arrays of derived types are multiply-added in d%y = a * d%y + c type equations: the RHS is put into an unnecessary temporary stack buffer and then copied to the LHS.

With large derived-type pointer arrays this both slows down performance and can use up stack memory. This will be rectified in a forthcoming version (version 12) of the compiler.

M. Plummer, K. Refson and Dr Matt Probert (York) have submitted a bid into the HECToR CSE Support call for a one-year project to introduce efficient band parallelization of CASTEP. Parallelization over bands will address both memory distribution (parallelizing the remaining CASTEP serial memory fraction) and performance by introducing a guaranteed new layer of parallelism, thus requiring the 3d-FFT only to scale over a fraction of the processors in any given calculation. If successful, this project will feed back to HPCx (and other high performance systems) and vastly improve both strong and weak scaling.

CRYSTAL: I.J. Bush has worked on two main areas of improvement this quarter.

1) Two important convergence acceleration techniques, the Anderson and Broyden methods, now have a distributed memory implementation. Previously they were very memory and I/O expensive as they require a large number of Fock/Kohn-Sham matrices to be held. The new implementation cuts down on the memory usage, and eliminates the need for I/O entirely.

For example, for a calculation on an Aluminium Fluoride surface both of the new implementations reduce the number of SCF cycles required from 25 to 16.

2) The start up of CRYSTAL has been optimised. It had been found that for large systems while the SCF is now very efficient the initialisation was taking longer than desirable. The start up time can itself be divided into 2 components

a) Initialisation that needs to be done once per calculation, i.e. only done once whatever the calculation
b) Initialisation that needs to be done once per geometry, i.e. may need to be done many times if the calculation is a geometry optimisation

For instance for a calculation on a large amorphous silica surface (~15,000 basis functions) it was found that a) was taking around 900 seconds, and b) 700 seconds. Work has reduced the time for b) to less than 10 seconds by refactoring portions of the code so that they are now suitable for parallelisation. Work is in progress on a), but should result in a comparable saving in time, and also a marked reduction in the memory footprint of the program

This work has been facilitated by two trips to Turin. At the second the work on the distributed memory version of CRYSTAL was presented at MSSC2007

VASP: K. D'Mellow has addressed some porting and numerical issues in VASP, and some user queries and issues that have arisen.

QUICKSTEP: Following a request for in-depth report from Ben Slater of the e05 (Materials Chemistry) consortium, some investigation of the 'Quickstep' code has been undertaken by K. Stratford. It was found that the current xlf version (10.01.0.4) on the main machine caused problems with incorrect code generation causing a number of Quickstep tests to fail. After investigation it appears that xlf 10.01.0.5 solves the problem. The Terascaling team has requested that the updated compiler is made available on main machine.

4.3.3 Computational Chemistry Codes

DL_POLY: DL_POLY_3 version 08 was finally released in September after 10 months of work by I.T. Todorov and I.J. Bush. In the last quarter, three release candidates (RC) have been tested by external users and feedback taken to improve quality and apply fixes where needed. A short list of important changes between version is provided at the end of the report.

RC1 has been used to benchmark performance and I/O speed on current cutting edge platforms; Cray XT3 SC/DC, IBM BG/L and HPCx. The performance wise order is Cray XT3 SC followed very closely by Cray XT3 DC, and IBM P5 in the middle between them and BG/L. It was found that for the most memory and MPI demanding part of the code (k-space Ewald) HPCx performed nearly as well as Cray XT3 SC which is much better than Cray XT3 DC followed closely by BG/L. The results were reported together with examples of two scientific studies carried out with DL_POLY_3 on HPCx at Parallel Computing 2007 in Jülich, Germany.

The I/O handling order is IBM P5, Cray XT3 DC, Cray XT3 SC and BG/L. It is worth mentioning that XT3 SC and DC have some fundamental problems with the traditional gather write criterion and only a type of a parallel write works for them. This is now being reported in a separate paper.

Here follows a short list of new features and important changes that have taken place from version 07 to 08:

- The Force-Shifted Coulombic (FSC) evaluation option is extended to cover the Fennell type of Ewald evaluation. This is optional and does not substitute the traditional FSC directive.
- Infrequent k-space Ewald evaluation option. This option has essentially the same effects (and the drawbacks) as the multiple timestep option in DL_POLY_2.
- "no strict" option for no strict global index contiguity check.
- "impact" option for setting impact energy and direction on a system particle specified by its global index.
- "regauss temperature" option to resample the temperature in the system.
- "spme sum" option introduced for backwards compatibility with DL_POLY_2. This has changed the meaning of pre-3.08 "ewald sum" option.
- NVT Andersen ensemble.

- NPT Langevin ensemble.
- NsT Langevin ensemble.
- The "minimise" option invokes a Conjugate Gradient Minimiser (CGM) for static structure minimisation.
- Important corrections to frozen particle contributions in the Ewald summation.
- Important corrections to MSD and particle displacement estimators.
- New local particle density is used to estimate array bounds dependent on density at start up time. It makes the use of "densvar" option less needed but not redundant.
- More stringent warning facility and safety checks.
- More stringent compliance of the code to the FORTRAN95 standards.

4.3.4 Physics Codes

GWST code: K. D'Mellow has continued to parallelise the GW-Spacetime code (GWST), and profiled further using a larger simulation set, including some large user simulations (Silicon surface and wire calculations). Development of the parallel version is performed in such a way that serial compilation and execution of the parallel code is possible throughout. The code now demonstrates correct I/O behaviour in parallel, and loop level parallelisation of key routines of the code are currently being addressed: specifically the data redistribution prior to and after various FFT stages within the code. Further porting issues have been fixed in sections of the supporting applications needed to preprocess GWST input datasets.

ang95: A.G. Sunderland has analysed the performance of the OpenMP-based inner region R-matrix code ang95 on HPCx. The results have been distributed to developers and users. The performance of the multi-threaded code on HPCx appears to have improved markedly of late and the speed of the code now increases steadily up to 16 threads (this is around 6 times faster than the serial version).

4.3.5 Engineering Codes

CFX: Following a user request CFX has been re-licensed for the system. CFX Version 11 is now installed and is available to all HPCx users until the end of the service in 2008.

STAR-CCM+: Following a user request STAR-CCM+ has been installed on HPCx.

4.3.6 Environmental Science Codes

OCCAM: F. Reid has written, tested and benchmarked a task farm harness for the OCCAM ocean modelling code so that the researchers at Proudman

Oceanographic Laboratory can run many instances of OCCAM simultaneously. This has subsequently been awarded a Gold Capability incentive. Writing the task farm harness involved altering various parts of the OCCAM code. The I/O in particular required special handling to ensure that I/O from each task was kept separate from other tasks. The researchers are now testing and running the new version of the code.

4.3.7 New Applications

STAR-CD: J. Hein has installed STAR-CD on the service machine, and has worked on associated complications due to license server issues.

4.3.8 Libraries

Global Arrays: Work is underway on timing of parallel libraries, Global Arrays (GA) and the Distributed Data Interface (DDI), running over LAPI and/or MPI-2, on HPCx. The measured bandwidths and latencies for key data access operations (GET, PUT, ACC) will help to predict the viability of candidate distributed data self-consistent field algorithms in the 'petascale' arena by parameterizing the associated performance models. The study is also covering XT4, BG/L, and Opteron/Myrinet clusters. Related work is underway on updates to DDI (DDI 3.2) incorporating fixes for LAPI, interfacing to MPI-2, and a new high-level suite of functions for manipulating whole distributed matrices ('DDA' routines).

Linear Solvers: A.G. Sunderland has analysed the performance of the latest ScaLapack eigensolvers for the R-Matrix codes on HPCx. The performance of multi-threaded eigensolvers has also been analysed and performance is being compared to other machines. This work was presented as an invited talk at the CCP2 Workshop on Mathematical and Computational Methods in R-matrix Theory at UCL in September. An article describing this work will be included in the published workshop proceedings.

RAMDISK I/O Performance: This research has not been progressed during this quarter due to a lack of suitable maintenance sessions for testing.

4.3.9 Tools

An Technical Report on HPCx usage of profiling tools, including the latest releases of Vampir/VampirTrace (v.5) and Paraver, has been published. This was followed up by a presentation on profiling tools to the HPCx User Group meeting in August. Further testing of the new Vampir and VampirTrace products is being undertaken with a view to purchasing a new license for HPCx.

Totalview: Totalview has been upgraded to version 8.2 on HPCx. This resolved some stability problems with the license server.

4.3.10 Applications on New Technologies

J. Hill has been investigating the HPC potential of the Cell Processor. This processor, found in the Playstation 3 console, could have huge impact in HPC due to the fact it has a total of nine processing cores in each chip and will be a relatively cheap commodity chip. A Petaflop supercomputer including the Cell (RoadRunner) has already been announced. However, the Cell technology is relatively new and programming it presents significant challenges. This work has investigated porting a very simple (but representative) code to the Cell processor, measuring the performance and detailing how to use the Accelerated Library Framework methodology of writing code for the Cell.

J. Hein has been preparing papers for publication on Fast Fourier Transforms and Blue Gene applications, both of which having a strong HPCx component.

4.3.11 Talks and Publications

Investigation of a set of real applications using Power5 hardware counters, Mark Bull, ScicomP13, IPP Garching, Munich, Germany, 16th-20th July 2007

Mixed Mode Programming on HPCx, Michal Piotrowski, ScicomP13, IPP Garching, Munich, Germany, 16th-20th July 2007

Fluent on HPCx, A.G. Sunderland, HPCx Technical Report HPCxTR0705

Chapel, Fortress and X10: novel languages for HPC, M. Weiland, HPCx Technical Report HPCxTR0706

DL_POLY_3 I/O Analysis, Alternatives and Future Strategies, I.T. Todorov and I.J. Bush, HPCx Technical Report HPCxTR0707

Parallel Diagonalization Methods for R-matrix calculations, CCP2 Workshop on Mathematical and Computational Methods in R-matrix Theory, University College London, UK, 14th-15th September 2007

Tools for Profiling Parallel Performance, A.G. Sunderland, HPCx User Group Meeting via Access Grid, 16th August 2007

DL_POLY_3: Parallel Performance and Large Scale Simulations, I.T. Todorov, W. Smith and I.J. Bush, Parallel Computing, FZ Jülich, Germany, 4th-7th September 2007

High Resolution Modelling of the Northwest European Shelf Seas using POLCOMS, J.T. Holt, R. Proctor and M. Ashworth, Capability Computing

Petascale Computing, M. Ashworth, UK Turbulence Consortium Workshop, University of Southampton, UK, 26th -27th September 2007

J. Hein has written an article for the upcoming edition of Capability Computing on the ScicomP event.

4.4 Software Engineering (Dr Stephen Booth)

4.4.1 Migration to HECToR

Significant effort has been invested this Quarter toward ensuring a straightforward migration of codes from HPCx to HECToR. This has resulted in a significant number of application codes being available to users at the start of the HECToR service.

4.4.2 In-depth software support

Grid middleware support

The Grid middle-ware support on HPCx has been updated this quarter. This has been driven by the expected requirements of the GENIUS project that requires the most up to date versions of Globus. We have been investigating both source code builds of globus-4.0.5 and the installation of binary packages from the Virtual Data Toolkit. The Virtual Data Toolkit is a repackaged binary release of a variety of grid software including globus. As VDT forms the basis of the NGS minimal software stack this potentially provides a route to greater NGS integration and easier maintenance. Unfortunately, although we successfully installed VDT 8.1 on HPCx, the level of AIX support from the VDT is still too rudimentary to provide any significant advantage over a from source build.

The HPCx contribution to the NGS was reviewed by the NGS in September. The overall use of HPCx compute time via the Globus interface remains very small. This is not too surprising as many of the potential benefits of computational grids only apply when there are a number of roughly equivalent target platforms to choose between. On the other hand the Globus tool-kit still provides the only tools for efficiently moving large data sets between different sites.

User administration software

Development of the SAFE system has continued this quarter. Though a number of additional features have been added to support the Hector service the changes relevant to the HPCx service have mostly been fine tuning the helpdesk system and general maintenance. Particular effort has been invested in the code documentation to make it easier for new staff to support this code.

4.5 Operations and Systems *(Mr Mike Brown)*

Reliability: There were no SEV1 incidents (ie, contractual failures) this quarter. The last failure attributable to IBM was in December 2006; the last attributable to the site was in September 2006. The reliability remains solid with only a few minor incidents.

Software upgrades: A1536 processor queue was made available.

Staffing: Colin Morey and Viliam Kalavsky joined the systems group at Daresbury. Iain Georgeson has tendered his resignation and will be leaving In October.

4.6 Staffing

AV	*Ju;y*	*August*	*September*
DL	4.4	4.5	4.4
EPCC	10.6	11.3	6.2
Total	15.0	15.8	10.6

Systems	5.9	5.4	6.0

5 Summary of Performance Metrics

Metric	*TSL*	*FSL*	*July*	*August*	*September*
Technology serviceability	80%	99.2%	100.0%	100.0%	100.0%
Technology MTBF (hours)	200	300	∞	∞	∞
Number of AV FTEs	7.5	10	15.0	15.8	10.6
Number of training days per month	22.5/12	30/12	13/7	13/8	13/9
Non in-depth queries resolved within 3 days	85%	97%	100.0%	96.5%	92.9%
Number of A&M FTEs	3.75	5.75	5.9	5.4	6.0
A&M serviceability	80%	99.6%	100.0%	100.0%	100.0%

Colour	*Meaning*
	Exceeds FSL
	Between TSL and FSL
	Below TSL

Note 1: The number of training days is reported as a running total since the start of the year.

Note 2: The above table includes the revised FSL targets for *training days* and *A&M serviceability*, which have been provisionally agreed with EPSRC.

Appendix A: Incident Severity Levels

SEV 1 — anything that comprises a FAILURE as defined in the contract with EPSRC.

SEV 2 — NON-FATAL incidents that typically cause immediate termination of a user application, but not the entire user service.

The service may be so degraded (or liable to collapse completely) that a controlled, but unplanned (and often very short-notice) shutdown is required or unplanned downtime subsequent to the next planned reload is necessary.

This category includes unrecovered disc errors where damage to filesystems may occur if the service was allowed to continue in operation; incidents when although the service can continue in operation in a degraded state until the next reload, downtime at less than 24 hours notice is required to fix or investigate the problem; and incidents whereby the throughput of user work is affected (typically by the unrecovered disabling of a portion of the system) even though no subsequent unplanned downtime results.

SEV 3 — NON-FATAL incidents that typically cause immediate termination of a user application, but the service is able to continue in operation until the next planned reload or re-configuration.

SEV 4 — NON-FATAL recoverable incidents that typically include the loss of a storage device, or a peripheral component, but the service is able to continue in operation largely unaffected, and typically the component may be replaced without any future loss of service.

Appendix B: Projects

B.1 Current Projects

EPSRC Projects

Code	*Class*	*Title*	*PI*
e01	1	UK Turbulence Consortium	Dr Gary Coleman
e05	1	Materials Chemistry using Terascaling Computing	Prof Richard Catlow
e08	2	Organic Solid State	Prof Sarah Price
e10	1	Reality Grid	Prof Peter Coveney
e19	1	Edinburgh Soft Matter and Statistical Physics Group	Prof Michael E Cates
e24	1	DEISA - Distributed European Infrastructure for Supercomputing Applications	Dr David Henty
e25	1	Turbulent vortex motion in stratified flows	Dr Gary Coleman
e26	1	Simulation of Radioprobing	Dr Charlie Laughton
e30	1	Metal/Oxide Interfaces at the Atomic Level	Dr Nora de Leeuw
e33	1	Engineering Functional Coatings	Prof Roger Smith
e34	1	Dissolution of Bioactive Phosphate Glasses	Dr N de Leeuw
e35	1	Non-adiabatic processes	Dr T Todorov
e36	1	Jets in Cross-Flow	Dr Y Yao
e37	1	LESUK_3	Prof J J McGuirk
e38	1	Viscoelastic deformation in 3D non-linear media	Prof Greg A Houseman
e39	1	The Supergen 5 biological fuel cells consortium	Prof FA Armstrong
e40	1	Computational Quantum Many-Body Theory	Prof R Needs
e41	1	Flow in Weapon Bays	Dr George N Barakos
e42	1	Computational Combustion for Engineering Applications	Prof K Luo
e44	1	Extreme Wave Loading on Offshore Wave Energy Devices	Dr Deborah Greaves
e45	1	Metals under extreme conditions	Prof Mike Gillan

e46	1	Advanced materials with complex architectures	Dr Paul Mummery
e47	1	Parallel stochastic analysis for geo-engineering	Dr Michael A. Hicks
e49	1	Integrated Programme of Research in Aeronautical Engineering	Prof Michael Leschziner
e50	1	Biological interface with materials	Prof John Harding
e53	1	Large-scale communication networks	Prof J M Pitts
e55	1	High-Reynolds-Number Near-Wall Flows	Prof Michael Leschziner
e57	1	Triplex DNA Structures	Dr Hooshang Nikjoo
e58	1	Quantum Simulations for Chemical Biology	Dr Carmen Domene
e59	1	Turbulence in breaking gravity waves	Prof Ian Castro
e60	1	Device Electronics Based on nanoWires and NanoTubes	Dr Merlyne M de Souza
e61	1	Enhancement of droplet concentrations in clouds	Dr Alan A M Gadian
e62	1	Low Voltage Defibrillation	Dr Vadim Biktashev
e63	1	UKAAC	Dr Nick Hills
e64	1	Jet noise from instability mode interactions	Prof Neil Sandham

PPARC Projects

Code	*Class*	*Title*	*PI*
p01	1	Atomic Physics and Astrophysics	Prof Alan Hibbert

NERC Projects

Code	*Class*	*Title*	*PI*
n01	1	Large-Scale Long-Term Ocean Circulation	Dr David Webb
n02	1	NCAS	Prof Alan J Thorpe
n03	1	Computational Mineral Physics Consortium	Dr John Brodholt
n04	1	Shelf Seas Consortium	Dr Roger Proctor

BBSRC Projects

Code	*Class*	*Title*	*PI*
b08	1	IntBioSim	Prof M S Sansom

CCLRC Projects

Code	*Class*	*Title*	*PI*
c01	1	Daresbury Laboratory Facilities Agreement Consortium	Dr Richard J Blake

Externally-funded Projects

Code	*Title*	*PI*
x01	HPC-Europa	Dr Judy Hardy

HPCx Projects

Code	*Title*	*PI*
z001	HPCx Support	Dr Alan Simpson
z002	Systems and Operations	Mr Mike Brown
z003	Test Project	Dr Denis Nicole
z004	HPCx Training	Dr David Henty
z05	Outreach Projects	Dr Richard Blake
z06	Application Porting	Dr David Henty
z07	Package Installation	Dr Mike Ashworth
z10	Globus	Dr Stephen P Booth

B.2 Former Projects

Code	*Class*	*Title*	*PI*
b01	2	Quantum Chemistry Studies of the Rusticyanin Protein Crystal	Prof Samar Hasnain
b03	1	Towards a virtual outer membrane	Prof Mark S Sansom
b04	1	Life sciences software development	Dr Jo L Dicks
b05	1	Virtual forced evolution of catalytic transition metal complexes	Dr Marcus Durrant

b06	2	Biomolecular computational chemistry	Prof Jonathan D Hirst
e02	1	Ab-initio simulation of covalently bonded materials	Dr Patrick Briddon
e03	1	Multi-photon, electron collisions and BEC HPC consortium	Prof Ken Taylor
e04	1	Chemreact Computing Consortium	Prof Jonathon Tennyson
e06	1	UK Car-Parrinello Consortium	Prof Paul Madden
e07	2	Turbulent Plasma Transport in Tokamaks	Dr Colin M Roach
e09	2	Molecular Properties and their Geometry	Dr Mark R Wilson
e11	1	Bond making and breaking at surfaces	Prof Sir David A King
e12	1	Parallel programs for the simulation of complex fluids	Dr Richard J Blake
e13	1	TeraGyroid project	Mr Mark Westwood
e14	1	Blade and Cavity Noise	Prof Neil Sandham
e15	2	CSAR/HPCx Collaboration	Dr Mike Pettipher
e16	1	Cardiac virtual tissues	Prof Arun V Holden
e17	1	Integrative Biology	Dr David Gavaghan
e18	1	DARP: Highly swept leading edge separations	Prof Michael A Leschziner
e20	1	UK Applied Aerodynamics Consortium	Dr Ken Badcock
e21	1	Intrinsic Parameter Fluctuations in Decananometer MOSFETs	Prof Asen M Asenov
e22	1	Preconditioners for finite element problems	Prof David J Silvester
e23	1	Exploitation of Switched Lightpaths for e-Science Applications	Prof Peter Clarke
e27	1	SPICE	Prof Peter V Coveney
e28	1	Towards the Dynome	Dr Jonathan W Essex
e29	1	Free-surface-piercing circular cylinders	Dr Eldad Avital
e31	1	Lateral Straining of Wall-Bounded Turbulence	Dr Gary N Coleman
e32	1	Rapid Prototyping of Usable Grid Middleware	Prof Peter V Coveney
e48	1	Organised structure in turbulent flows	Prof Sergei Chernyshenko

e51	1	Super-computing data mining	Dr Mike Pettipher
e52	1	Spacecraft force modelling	Dr M Ziebart
e54	1	Free surface simulation of waves overtopping during storms	Dr D M Ingram
e56	1	Infectious disease threats	Dr Iain Barrass
b02	1	Modelling enzyme catalysis	Dr Adrian J Mulholland
z09		HECToR Benchmarking	Dr Edward Smyth
x02		OHM Ltd	Dr Lucy MacGregor
x03		IBM	Mr Derrick J Byford
n05	2	Non-linear Wave-particle Instabilities in Plasmas	Dr Mervyn Freeman