HPCx Quarterly Report
April – June 2006

1 Introduction

This report covers the period from 1 April 2006 at 0800 to 1 July 2006 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

- The system continues to be reliable, with only four failures this quarter. Preparations for the upgrade to Phase 3 are still on track and IBM has carried out some preliminary work.
- Although there was a sharp fall in utilisation in April, it is recovering. Utilisation of the capacity region is very high; utilisation of the reserved development region continues to increase.
- Thanks to collaboration between IBM and HPCx staff, the Simultaneous Multi-Threading (SMT) facility of the POWER5 processors is now available to users. A technical report on this has been published.
- Following authorisation by EPSRC, twelve CSAR projects were moved to HPCx in the period leading up to the closure of the CSAR service. A joint workshop with CSAR was held in Manchester in May to help CSAR users make the transition.
- There are now 53 projects on HPCx, including the former CSAR projects, with another approved by EPSRC for access. This leaves one spare place within the new maximum of 55.
- Six courses were run this quarter over a total of 20 days. This now puts our training activity well ahead of the targets.

- We are continuing to progress well against the Key Objectives in the Annual Plan. Another capability incentive was awarded. Six out of ten Technical Reports planned for the year have been completed. We met all the helpdesk targets and performance metrics continue to be good.

- David Henty's talk at the Edinburgh International Science Festival, *Supercomputing: Rise of the Machines*, which strongly featured HPCx, was well attended.

- The latest edition of *Capability Computing* is being mailed to nearly 4,000 people and is available online. A meeting of the User Group was held via Access Grid, with 11 users taking part; topics included SMT.

- This year's annual conference, *Moving Science Forward*, will take place in the National e-Science Centre in Edinburgh on 4 October and a number of speakers have already been confirmed. A workshop on Materials Modelling will be held on the preceding day, also in Edinburgh.

- The number of packages and libraries supported is now approaching 60.

- We have enhanced our support for Grid middleware, allowing users of Globus to run MPI programs as metacomputing jobs across multiple sites --- a Technical Report has been published discussing this.

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:

	April	*May*	*June*
Incidents	10	14	14
Failures	1	3	0

The following tables give more details on the attribution of the failures:

April

Failure	*Site*	*IBM*	*External*	*Reason*
06.036	0%	100%	0%	Maintenance session over-run

May

Failure	*Site*	*IBM*	*External*	*Reason*
06.045	0%	0%	100%	External network failure
06.046	0%	0%	100%	External network failure
06.051	0%	100%	0%	Maintenance session overrun

June

None

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*	*April*	*May*	*June*	*Quarterly*
IBM	Failures	1	1	0	2
	MTBF	732	732	∞	1098.0
	Serviceability	99.7%	99.3%	100.0%	99.7%
Site	Failures	0	0	0	0
	MTBF	∞	∞	∞	∞
	Serviceability	100.0%	100.0%	100.0%	100.0%
External	Failures	0	2	0	2
	MTBF	∞	366	∞	1098.0
	Serviceability	100.0%	98.4%	100.0%	99.5%
Total	Failures	1	3	0	4
	MTBF	732	244	∞	549.0
	Serviceability	99.7%	97.7%	100.0%	99.1%

3.2 Capability Utilisation

The monthly utilisation for the capability region of the main service (the production region on Phase 2) is shown in the graph below. A sharp drop in April has been followed by some recovery.

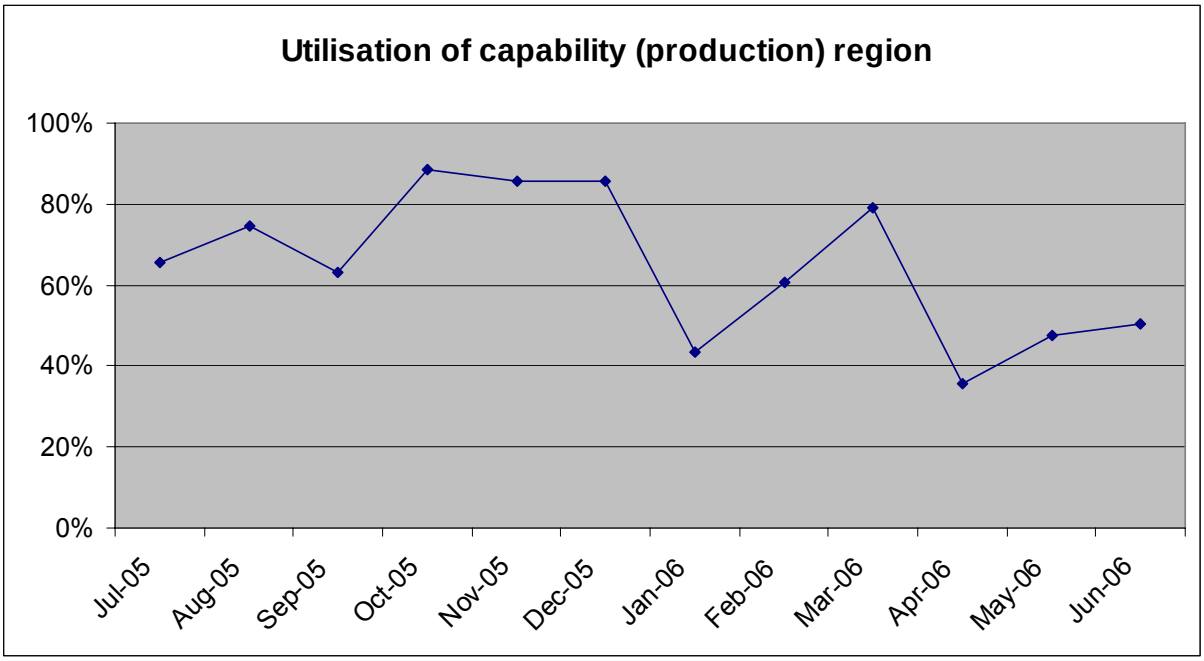

3.3 Capacity Planning

Predicted Utilisation

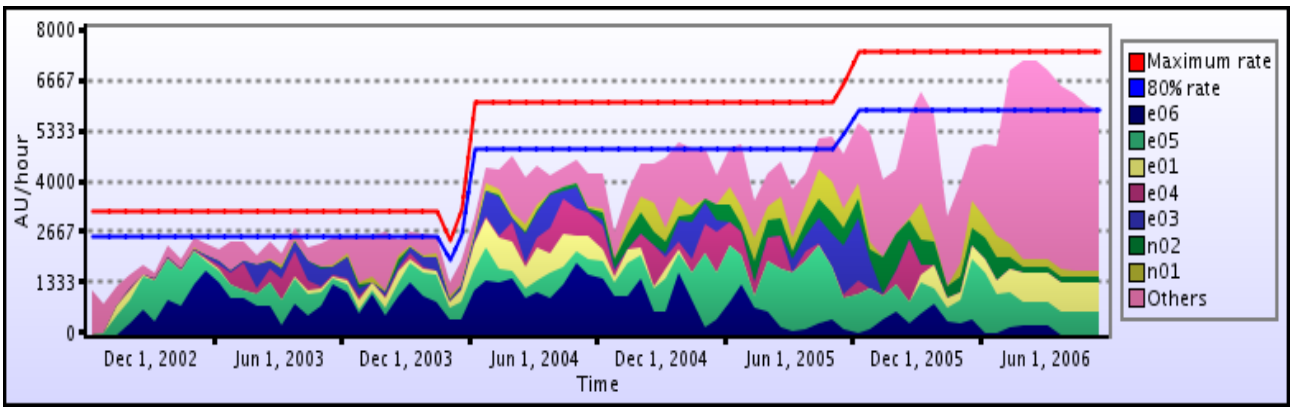

The graph above shows the utilisation since the start of the project and the projected utilisation (on the main service) until December 2006. The scale on the *y*-axis is AUs per hour, where at peak Phase 2A can deliver 7395 AUs per hour (the upper red line in the graph). The lower line (in blue) corresponds 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 there is likely to be some spare capacity later this year and next, especially when Phase 3 comes on stream.

Numbers of Research Consortia

There are currently 53 research consortia on HPCx. They include twelve projects which have been moved from CSAR as a result of the closure of that service. Another project has been approved for access by EPSRC.

In addition, there is one active externally funded project.

3.4 CPU Usage by Job Size

Main service

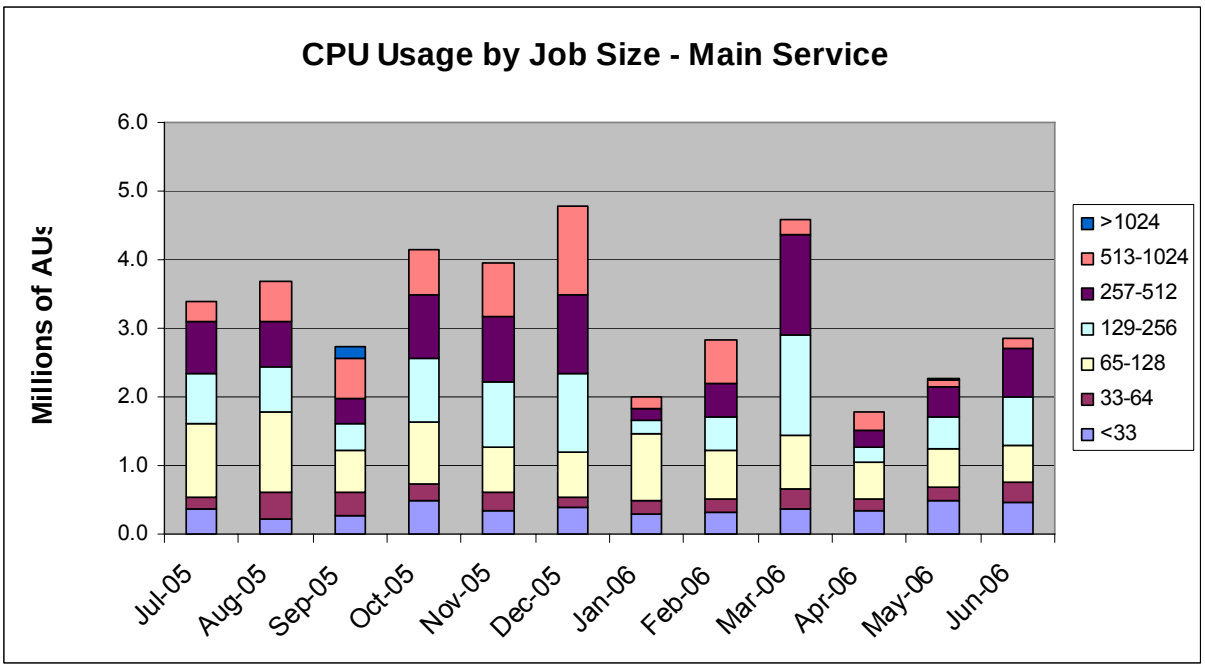

Development Service

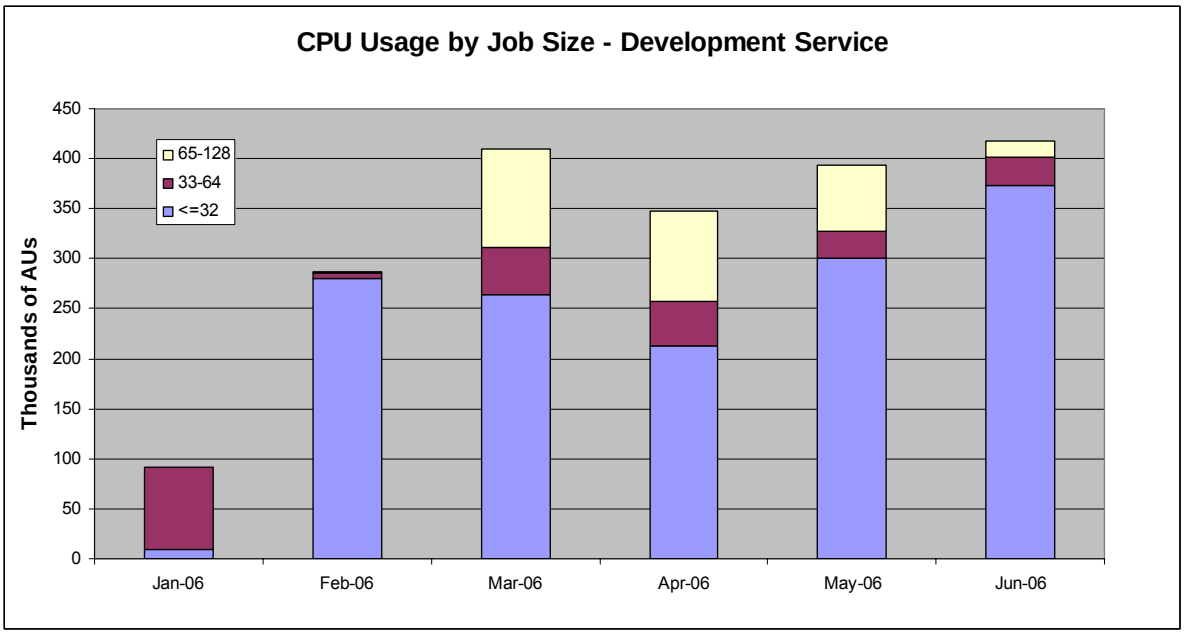

3.5 AU Usage by Consortium

Main Service

Consortium	*April*	*May*	*June*	*Quarterly*	*%age*
e01	192431	238732	258438	689601	8.3%
e03	58			58	0.0%
e05	238486	918088	954627	2111201	25.4%
e06	295643	252201	59826	607671	7.3%
e07	25009	177741		202750	2.4%
e08	47531	118502	60626	226659	2.7%
e10		646	21231	21876	0.3%
e11		5347	43163	48510	0.6%
e14	280635	42412	347590	670638	8.1%
e15	543		6	549	0.0%
e17	23766	13275	129970	167012	2.0%
e18	22453	0	0	22453	0.3%
e19	0	2	350	352	0.0%
e20	75153	68144	259	143557	1.7%
e21			308	308	0.0%
e23			12	12	0.0%
e24	22608	52	846	23506	0.3%
e25	26762	278	3497	30536	0.4%
e26	5		3555	3560	0.0%
e27	75077	59297		134374	1.6%
e28	96908			96908	1.2%
e31	29914	48603	7180	85698	1.0%
e32	148874	110734	2842	262450	3.2%
e33	9370	39625	35084	84079	1.0%
e35	4042	12579	1032	17653	0.2%
e36	3080	84407	22516	110003	1.3%
e37	41742	5911	108415	156068	1.9%
e40	49623	16444	19067	85134	1.0%
e41			0	0	0.0%
e49			1	1	0.0%
e50			147	147	0.0%
EPSRC Total	1709714	2213022	2080588	6003324	72.3%

n01	0	412772	350280	763052	9.2%
n02	129670	182354	167144	479168	5.8%
n03	75092	37519	49277	161888	1.9%
n04	63446	16538	134679	214662	2.6%
NERC Total	268208	649182	701380	1618771	19.5%

p01	24176	7163	42091	73429	0.9%
PPARC Total	24176	7163	42091	73429	0.9%

c01	103608	42989	202911	349509	4.2%
CCLRC Total	103608	42989	202911	349509	4.2%

b08	58810	47250		106059	1.3%
BBSRC Total	58810	47250	0	106059	1.3%

x01	19439	9996	14048	43483	0.5%
x03	47654			47654	0.6%
External Total	67093	9996	14048	91137	1.1%

z001	11193	8999	25142	45335	0.5%
z002	234	3138	26	3398	0.0%
z004		15	2835	2850	0.0%
z05	114	1686		1800	0.0%
z06	9719	7		9726	0.1%
HPCx Total	21260	13845	28004	63109	0.8%

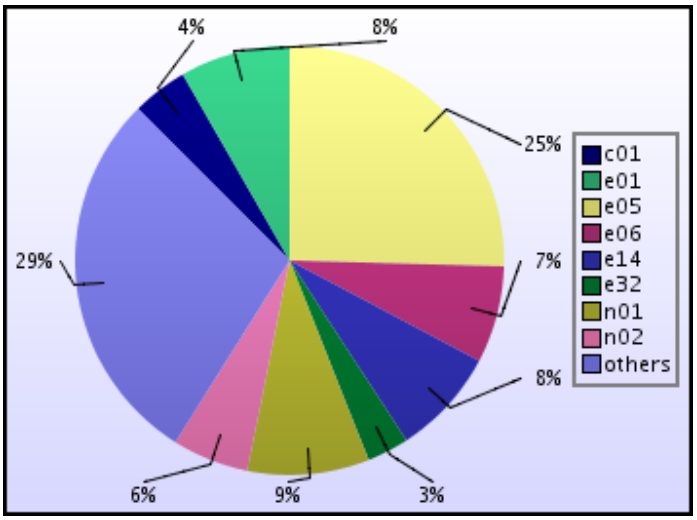

Development service

Consortium	*April*	*May*	*June*	*Quarterly*	*%age*
n01	9676	40508	16546	66731	5.8%
n02	135253	151401	207872	494526	42.6%
n03	203049	201217	194547	598813	51.6%
NERC Total	347979	393127	418964	1160070	100.0%

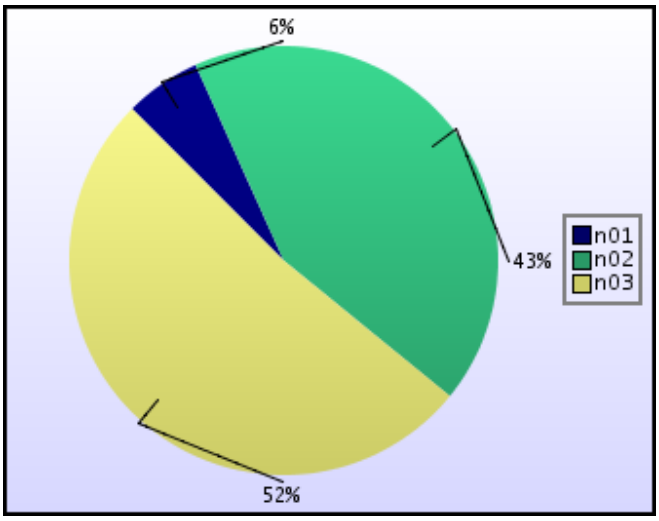

3.5.1 Discounts

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

Consortium	*AUs Used*	*AUs Charged*	*Discount*
c01	351640	349508	2131
e05	2131120	2111201	19918
e28	103298	96908	6390
e32	340628	262449	78178
e36	112591	110002	2588
e40	85232	85133	98

3.6 Helpdesk

3.6.1 Classifications

Category	*Number*	*% of all*
Administrative	161	49.8
Technical	144	44.6
In-depth	15	4.6
PMR	3	0.9
TOTAL	323	100.0

Service Area	*Number*	*% of all*
Phase 1/2 platforms	283	87.6
Website	15	4.6
Other/general	25	7.7
TOTAL	323	100.0

3.6.2 Performance

All non-indepth queries	*Number*	*%*	*Target*
Finished within 24 Hours	262	85.9	75%
Finished within 72 Hours	303	99.3	97%
Finished after 72 Hours	2	0.7	

Administrative queries	*Number*	*%*	*Target*
Finished within 48 Hours	161	100.0	97%
Finished after 48 Hours	0	0.0	

3.6.3 Experts Handling Queries

Expert	*Admin*	*Technical*	*In-Depth*	*PMR*
epcc.ed.ac.uk	122	65	6	0
dl.ac.uk	4	27	4	3
Sysadm	34	51	5	0
Other people	1	1	0	0

3.7 Service Quality Tokens

Date	*Person*	*Value*	*Comment*	*Status*
Jul 1, 2006 1:54:28 AM	*Dr Glenn D Carver*	* *		
Jun 6, 2006 2:59:51 PM	*Dr Jeffrey M Chagnon*	* * * *		

4 Support

4.1 Applications Support *(Dr David Henty)*

4.1.1 Documentation

Since the Phase 2A upgrade, the stability of the service has meant that documentation has remained fairly stable with only minor updates and modifications required. However, the introduction of the new Simultaneous Multithreading (SMT) facility and the arrival of a large number of new users from the CSAR service lead to a number of significant changes. We have also maintained contact with users via six user mailings in Q2.

The HPCx system has in principle supported the new SMT feature since the delivery of the new Power5 technology with Phase 2A last year. However, it was only in Q2 this year that the system software was sufficiently mature to allow SMT to be enabled for general users. Information on SMT has been added to the User Guide and as an FAQ entry. The documentation has links to the associated technical report, and a discussion of SMT was included in a talk at the recent User Group (see below for more details on these).

We received feedback from the helpdesk that there was confusion among some new users regarding the setup of the LoadLeveler batch system. We have added a section to the User Guide giving a brief overview of the allowable job sizes and durations since these are substantially different from CSAR. We also expect an increase in usage of OpenMP with the new users, so we have added an FAQ entry detailing how to run such jobs on a capability service such as HPCx.

4.1.2 Technical Reports

Three reports were planned for Q2 in the following areas:

a) Applications Performance on Phase 2A
b) Achieving Capability Incentives for HPCx Applications
c) HPC Software Survey

Report (a) was already delivered in Q1 as HPCxTR0602. We have also produced the following three reports this quarter:

- **HPCxTR0604**: *An Investigation of Simultaneous Multithreading on HPCx*, A. Gray et al.

- **HPCxTR0605**: *Terascale materials modelling on high performance system HPCx (J Mater Chem 16, 2006, p1885)*, M. Plummer, J. Hein, M.F. Guest et al.
- **HPCxTR0606**: *Grid Metacomputing support on HPCx*, S. Booth.

Report **04** was originally planned for Q3 but was produced this quarter as the SMT facility became available on HPCx earlier than anticipated (work by the HPCx systems team meant it could be deployed ahead of IBM's own schedule). Report **05** was originally planned for Q4 as it was expected that its publication in the Journal of Materials Chemistry would cause some delay. However, the publication process was very rapid and hence this report was made available to users ahead of schedule.

Report **06** was written in response to requests from certain user groups, most notably Peter Coveney's RealityGrid project, on how to run metacomputing jobs across multiple supercomputers including HPCx. A number of technical challenges were overcome by the software engineering team in order to make metacomputing on HPCx possible. A technical report was produced as the solution was sufficiently general to be of interest to other users.

Although there has been some adjustment of the schedule in response to changing circumstances, we are well on target with respect to the Annual Plan. A total of five reports were due by the end of Q2: we have actually produced six reports so far, with five of the titles coming directly from the original plan.

4.1.3 Training

In Q2 of 2006 we ran the following six courses, all at EPCC.

- **18 – 20 April:** *Fundamental Concepts of HPC;*
- **25 – 28 April:** *Practical Software Development;*
- **2 – 4 May:** *Shared-Memory Programming using OpenMP;*
- **9 – 11 May:** *Message-Passing Programming using MPI;*
- **24 – 26 May:** *Parallel Decomposition;*
- **30 May – 2 June:** *Applied Numerical Algorithms.*

Statistics are summarised below alongside annual targets (where appropriate):

Metric	Total	Target
Course days	20	25
Different course titles	6	8
Different locations	1	4
Student-days for HPCx users	217	
Student-days for HPCx staff	21	
Student-days available for HPCx	420	

As explained in the Q1 report, the start of the training schedule for this year was delayed until mid-April in order to align the HPCx programme with that of the CSAR service. The courses run this quarter also formed part of the HEC Training Centre programme and there is substantial leverage between the two training schedules (though note that HEC students are omitted from the statistics above). In the coming quarter we plan to run courses at different locations including Queen's University Belfast and Daresbury Laboratory.

4.1.4 Workshops and Conferences

The one-day workshop *Porting Codes from CSAR to HPCx* was held on 25 May at Manchester Computing, and was organised by HPCx in close collaboration with CSAR. The majority of material was presented by HPCx, but a talk on the differences between the SGI Altix and HPCx was given by CSAR staff. Experts from both services were available for all the practical sessions. The workshop was very successful and was attended by 12 users. For copies of all the talks see http://www.hpcx.ac.uk/support/training/HPCxCSARPorting.html. The second applications-based workshop will cover Materials Modelling and will be held in Edinburgh on October 3rd, immediately before the Annual Seminar.

This year's conference, *The Fourth HPCx Annual Seminar: Moving Science Forward*, will be held in Edinburgh on October 4th. We have obtained support from the National e-Science Centre who will host the event and contribute to speaker expenses. We already have a number of confirmed speakers and have started active publicity for the event. For registration and further details see http://www.hpcx.ac.uk/about/events/annual2006/.

4.1.5 User Group

As planned, we held the first User Group of 2006 over Access Grid. It took place on 28 June and was attended by 11 users at five different sites including Manchester, Cambridge and Southampton; representatives from the HPCx teams were also present. The timing was chosen to be useful for new users coming from CSAR. There were talks on the performance of the Phase 2A system including the use of SMT, and an overview of how to make best use of the HPCx tape archive. The format was very successful and we plan to continue to use Access Grid as an option for future User Groups. For copies of the presentations see http://www.hpcx.ac.uk/about/events/user_group_2006.html.

4.1.6 Newsletter

Production was slightly postponed as we had rather too much contributed material. The need to go back to the authors to approve changes to articles introduced a delay that we had not planned for. However, the newsletter is now completed and available online at http://www.hpcx.ac.uk/about/newsletter/. Hard

copies have recently been sent out to almost 4000 addresses including all HPCx users.

4.1.7 Packages

The number of pieces of software supported under the package mechanism is now approaching 60. Recent additions include the NCAR Graphics Library that was installed in response to a helpdesk request from a new user from CSAR.

4.2 Outreach Activities *(Dr Richard Blake)*

4.2.1 Lifesciences

A Report on the HPCx-IBM Lifesciences has been drafted and we are awaiting final input from a number of the more recent projects – DYNOME, Integrative Biology and Retina Modeling prior to finalising the report in Q3.

The Computational Science and Engineering Department held a meeting with BBSRC to discuss computational activities from the desktop to HECToR with programme managers and a number of areas for progressing Collaborative Computational Projects, HEC Consortia and distributed computing support activities with the academic community are being explored.

One measure of success of the HPCx-IBM Lifesciences programme is that the BBSRC involvement in High End Computing has grown from the 1% of usage on HPCx to a planned 5% allocation on HECToR.

4.2.2 Public Awareness

A public talk, *Supercomputing: Rise of the Machines* was given by David Henty on Sunday 9 April as part of the Edinburgh Science Festival. It was an event in the *Cutting Edge* theme of the *Big Ideas* programme. Other events in this theme included a presentation on *Cloning and Drug Discovery* from the leader of the team that produced Dolly the sheep, and a talk on *Pioneering Space Tourism* from the president of Virgin Galactic.

The talk was given to a general public audience and described how supercomputers are built and programmed, and the grand challenge science problems that they can tackle. The HPCx system was highlighted, and a number of applications described including results from the RealityGrid project and fusion simulations performed by the group at Culham. The event was extremely successful with an audience of 100 people ranging in age from school pupils to pensioners. For further details see the recent edition of Capability Computing.

4.2.3 Industrial Outreach

OHM Ltd bought significant CPU resources on HPCx last year to perform very large simulations for oil exploration. They have another code which was recently parallelised using OpenMP, and this has now been successfully ported to HPCx by EPCC's George Beckett despite a number of problems arising from moving from a Linux environment to AIX.

The aim was to evaluate performance and scaling on POWER5. The work was extremely successful with very good serial performance and a speedup of 6.5 achieved on 8 threads. This parallel efficiency of over 80% is an excellent result for a full-scale industrial code.

HPCx offers significantly more memory than the systems that OHM use on a daily basis and would allow them to run much larger problems (which we would expect to exhibit even better scaling than currently observed). We understand that solving large problems on a rapid timescale is a key business objective for OHM, and we are following up the possibility of selling them more time on HPCx.

CCLRC has pursued its customer contact programme in Q2 having meetings with AZ Mereside and Scienomics. A further series of presentations at British Nuclear Group, Nexia Solutions, Pilkingtons, AWE and Solvay Chemicals, has been organised for Q3. Presentations at these meetings explore hardware and software requirements on both capacity and capability systems. The major interest at the moment is in exploiting capacity systems where cost, network access, security and licensing of commercial software are the major issues. HPCx, as a capability system, finds it difficult, unsurprisingly, to compete on cost grounds in these high-throughput design applications.

For capability applications business appears to find it difficult to make the investment case to demonstrate the costs and benefits. CCLRC has recently established a ‘Proof of Concept’ fund to support its knowledge transfer activities, and projects which will support the internal business case for high-end and large-scale capability computing are being explored with our various industrial partners.

4.3 Terascaling Applications *(Dr Martyn Guest)*

4.3.1 Key Objectives

Progress against key objectives:

Capability Incentives
The following codes have obtained capability incentives this quarter:

- MERLIN_3D (Bronze)

SMT testing
Simultaneous Multithreading (SMT) has now been made available to users and can be enabled for individual jobs using simple LoadLeveler commands. The Technical Report comparing the performance of a number of application codes with and without SMT has been made available as HPCxTR0604.

4.3.2 Computational Materials Codes

CASTEP: As part of the plan for upgrading CASTEP discussed in the previous report, the lattice vibrations (phonon) part of the code has been adapted for use with the internal (intelligent) task-farming procedure. Martin Plummer has been discussing further extensions of the internal task-farming procedure with the CASTEP Developers Group (CDG) and the code has also been successfully adapted for use with the 'dumb' task-farming script (A.G. Sunderland) in a similar manner to that reported previously for DL_POLY_2. The main focus in the final part of the quarter is the investigation of the blocking of bands in 3d-FFTs to improve general performance. This will be completed in the next quarter. Various bug-fixes have been made to the released executables and discussion with the CDG (and testing) of recent changes made to basic routines in the top-level experimental version of the code is also underway.

Negotiations have also begun over setting up a 'fair' benchmark test case to compare (in confidence) the performance of CASTEP and VASP. A prescription for setting up such a case was put forward by the CDG (Keith Refson). Dario Alfè, who is an expert VASP user, has suggested some modifications. The CDG (Keith Refson) is studying these but wishes to defer the actual tests until completion of the current program of optimisations, which include optimisation of the gamma-point code and real-space pseudo-potentials.

Martin Plummer and Kenton D'Mellow attended the half-yearly meeting of the Materials Chemistry Consortium on behalf of HPCx.

VASP: Kenton D'Mellow has continued to investigate and improve the performance of VASP on HPCx on Phase 2A, focussing particularly on SMT

performance and MPI_Alltoallv. This work is currently being written up as a technical report.

4.3.3 Computational Chemistry Codes

DL_POLY: DL_POLY work this quarter has focused on enabling users to run larger scale calculations than previously possible.

Ilian Todorov and Ian Bush have worked to lessen the memory footprint of DL_POLY 3 (the distributed data version of the code). To help the user perform larger runs, the code has also been improved in two ways. The first is simply to provide a better estimate of the required resources for a given run. This helps users to estimate better both how large a system they can run, and also how many processors are required. The second is to improve the tools that generate the input for DL_POLY 3. The initial conditions of many MD simulations are such that they are simply a unit cell replicated many times in space, or at least small perturbations of such a structure. The code has been enhanced to help the user with such geometries by automatically producing them as required.

Some work has also been performed on examining how SMT affects DL_POLY. The results suggest that SMT improves the performance of the code somewhat at low processor counts, but at high processor numbers the improvement is negligible. The improvements do seem to compensate for the performance degradation observed on moving from Phase 2 to Phase 2A.

Some performance optimisation work has continued, especially an attempt to remove some of the square root evaluations in the code. However at present this is not noticeably improving the performance of the code. The reason for this is not clear.

GAMESS-UK: Recent large scale electronic structure calculations have shown a number of areas that need improvement in quantum chemistry codes: The time to solution, especially the integrals, grows as an unacceptable power of the system size (N^2 or worse), the memory requirements are excessive and the convergence for many systems is too slow. Members of the Terascaling team at DL are beginning to address these issues. In this quarter

1. An algorithm for the coulomb term that should scale approximately linearly with system size has been devised. This is based upon the KWIK algorithm of Gill et al, but with a new form of the long range component derived from the Ewald sum work of Ian Bush and Bill Smith in DL_POLY. The extension to deal with charge distribution. Implementation of this algorithm has been begun, although it is too early to give results.

2. A very initial examination of the recent method of Kohalmi et al, Phys Rev Let, 95, 013002, (2005) shows good promise both for the improvement of the convergence and eventually for moving to O(N) in both memory and time for

the linear algebra section of the SCF procedure by adopting a diagonalisationless, direct minimisation procedure.

3. An algorithm for a fully data distributed Fock / Kohn-Sham matrix build has been devised which fits well with 1) above.

During the last quarter the GAMESS-UK port to Blue Gene was released. Currently this is the MPI parallel implementation only. A study is underway to determine how the performance of the code compares against HPCx. Looking further ahead we have also been exploring transition state structure optimisations of complex catalytic systems using DFT. We have found that optimisation techniques well-suited to small systems can develop instabilities when applied to much larger (and more realistic) systems. The exact nature of this instability and possible solutions to the problem are being investigated.

GAMESS-US: Throughout the quarter, Graham Fletcher has been testing and benchmarking the quantum chemistry package GAMESS (US) on HPCx. So far work has principally focussed on key electronic structure functionality, such as multi-configurational self-consistent field (MCSCF), configuration interaction (CI), open-shell Z-averaged second order perturbation theory (ZAPT(2)), and effective fragment potentials (EFP), for which there is considerable interest within the user community. In addition, these efforts build toward a planned software exchange program with the developers of GAMESS (US) at Iowa State University (USA), with the aim of furthering advanced capabilities in the area of chemistry on HPCx.

SIESTA: J
Joachim Hein and Jon Hill have now installed SIESTA 2.0 on the service. Joachim and Andrew Sunderland have also analysed the source of SIESTA 2.0, to check for scope for future improvement. Finally, Joachim has started a joint collaboration with the code developers to analyse and improve the performance and scalability of this new release of SIESTA, with HPCx user supplied benchmarks.

4.3.4 Physics Codes

CENTORI: Joachim Hein has written, tested and benchmarked a dedicated Fourier transformation routine allowing for a 3-Dimensional parallelisation of this application. He has analysed the latest version of the application and produced a detailed strategy on how to implement the 3-Dimensional parallelisation. The required subroutines for the exchange of data between the processing nodes and for the flexible set-up to enforce boundary conditions have been implemented. Detailed test-routines have also been developed. The process has been documented in a report to the consortium (J. Hein, Using the parallel FFT set-up for CENTORI/CADENCE, June 2006). Finally he has liaised with Culham over their joint future PhD student.

Ludwig: Kevin Stratford has continued to provide support for the e19 (Edinburgh Soft Matter) project, and in particular, for the lattice Boltzmann code, Ludwig.

4.3.5 Engineering Codes

UKAAC: Two presentations were contributed to the UKAAC Consortium Conference on the 6th and 7th April 2006 at the University of Loughborough (see below). Two papers with contributions from HPCx staff have been submitted for publication in a Special Edition of the Aero Journal (see below).

Fluent: Version 6.3 of Fluent has been installed on HPCx. Among other features this version has improvements to the input/output and partitioning phases. Andrew Sunderland has been working with Kris Midgely at Fluent in order to build a benchmark dataset that reflects the size of problem (~100 million cells) that is intended to be solved on HPCx by users (existing Fluent benchmark datasets were far too small). The Fluent solver is now seen to scale well up to 512 processors and is therefore in line for a silver capability incentive award.

4.3.6 Life Sciences

Amber: Lorna Smith has installed the latest version of Amber on HPCx (Amber 9).

Gaussian: Lorna Smith and Judy Hardy carried out a survey of user requirements and experiences for Gaussian users (L. Smith, J. Hardy, Gaussian on HPCx, 2006).

Integrative Biology: Kevin Stratford has taken over support for the Integrated Biology (e17) project. He has investigated the performance of Simon Scarle's heart code, and provided a report to the consortia (K. Stratford, Simon Scarle's Heart Report, June 2006). The code investigates electrical activity in the heart and is written in C++ and MPI. Particular recommendations include changing the 1D decomposition to 3D and utilising a task farm harness. In addition, Kevin has been working with the users of the CARP (Cardiac Arrhythmia Research Package) code, providing support and upgrades for various libraries (PETSc, HYPRE, etc).

4.3.7 New Applications

ClockModel: Fiona Reid has been working with Professor Andrew Millar to port and test his code which models the biological clock of plants. Professor Millar has submitted a proposal for time on HPCx to BBSRC. Fiona has installed the code on HPCx and Blue Gene. The latest version of the code is written in C, removing the dependency on MATLAB. The code exhibited a number of

bugs/features which required interaction with the developers in order to resolve. The code has also tested with the task farm harness on HPCx. Fiona Reid has met with A Millar and S Setiawa to discuss future development with respect to improving the parallelisation.

4.3.8 Environmental Codes

OCCAM: A serial version of HDF5 has been installed so that users do not have to compile serial programs with the MPI compiler to simply use HDF5.

4.3.9 Libraries

Eigensolvers: Andrew Sunderland has been conducting performance tests of the beta version of the proposed ScaLAPACK MRRR-based algorithm PDSYEVR on a range of matrix characteristics. The code performed very well for certain matrices but experienced load-balancing problems on others. A report detailing the initial findings of this code has been sent to Christian Voemel (Berkeley) and the ScaLAPACK code developers who have subsequently provided an updated version of the code. This new code performs much better across the range of matrices and the report is currently being updated with the new data. At present, the report is internal, but will soon form the basis of an HPCx Technical Report for users due next quarter.

A.G. Sunderland is now looking at the application of the new parallel eigensolvers to suitable codes on HPCx. Three have already been identified: VASP, SIESTA and KPPW (multi-band k.p. code for quantum dot analysis). Initial benchmarking of the code for VASP matrices has commenced and results so far look promising.

4.3.10 Tools

TotalView: The TotalView memory debugging problems on HPCx have been fixed by application of a PTF from IBM.

DDT: The DDT debugger has been upgraded to version 1.9.2. Version 1.10 should be available within the next few weeks. The upgrade fixes a number of problems which had been reported to Allinea Software and includes memory debugging. When version 1.10 is installed it is planned to carry out an update of the comparison with TotalView (which was first reported in HPCxTR0506) and issue an updated Technical Report in Q3.

Java: Java 1.5 has now been released for AIX. This provides many additional concurrency features including barriers, locks and new spawning mechanisms. Lorna. Smith, Mark Bull and Paul Keir are investigating the potential benefits of these developments for parallel Java codes.

Hpmcount: Mark Bull is currently investigating hpmcount on POWER5, as many of the counters have changed from POWER4. This is to allow a report on the single node efficiency of application codes on the POWER5 system.

4.3.11 Talks and Publications

1. “Trends in High Performance Computing”, A.G. Sunderland and M. Ashworth, UKAAC Consortium Conference, 6th-7th April 2006, University of Loughborough.

2. “Historical Perspective on Aerodynamics and Supercomputing”, D.R. Emerson, UKAAC Consortium Conference, 6th-7th April 2006, University of Loughborough.

3. “Performance Benefits on HPCx Phase 2A: Power5 Chips and Simultaneous Multithreading”, Alan Gray, HPCx User Group Meeting, 28th June 2006.

4. “Use of the Tape Archive on HPCx”, A.G. Sunderland, HPCx User Group Meeting, 28th June 2006.

5. “HPC, HPCx and Blue Gene”, L. Smith, GSK / Edinburgh Bioinformatics Workshop, 20/21 June 2006.

6. “Novel Computing in Edinburgh”, J. Hein, The Irish Association for HPC, Cork/Ireland, June 2006.

7. "An Investigation of Simultaneous Multithreading on HPCx", A. Gray, J. Hein, M. Plummer, A.G. Sunderland, L. Smith, A. Simpson, A. Trew, HPCx Technical Report HPCxTR0604

8. “Simon Scarle’s Heart Report”, K. Stratford, internal report to consortia, June 2006.

9. “Using the parallel FFT set-up for CENTORI/CADENCE”, J. Hein, report to consortia, June 2006.

10. “Application of Parallel Rotor Simulation Tools on HPCx”, C.B. Allen, A.G. Sunderland, R. Johnstone, Aero Journal, Special Edition (submitted)

11. “High Performance Computing and Computational Aerodynamics in the UK”, D.R. Emerson, K. Badcock, A.G. Sunderland, M. Ashworth, Aero Journal, Special Edition (submitted)

4.4 Software Engineering *(Dr Stephen Booth)*

4.4.1 Future look and strategic HPC support

HPC user requirements survey
We made the Grid requirements survey available to users at the end of March and have advertised its existence several times in HPCx user mailings. Unfortunately there has been very little response from the users so far, possibly reflecting the impression we have from user support requests that the Grid is only important for a subset of the HPCx consortia. From the responses we have received so far the only noticeable trends are:

- DataGrid is important and needs to be fast/efficient and easy to use.
- Current users are not expecting to access databases over the Grid.
- There is no clear winner in terms of required middleware, though Globus 4 is probably going to be important.

4.4.2 Low-level investigations and analysis

We have been investigating the capabilities of the IBM C++ compiler and exploring the performance impact of various C++ language constructs using this compiler. Though the majority of HPC codes are still written in Fortran, the proportion of C++ codes appears to be increasing so a good understanding of the performance issues associated with this language is important.

Jon Hill, Adrian Jackson and John Spray have begun to investigate the HPC scientific potential of the new Cell processor. This processor will be in the forthcoming PlayStation 3 console, but could also have an impact for HPC systems due to the fact it has a total of nine processing cores in each chip giving it a potentially high peak performance but will be a relatively cheap commodity chip.

4.4.3 In-depth software support

We have installed the **boost** library on HPCx. This is a commonly used library of C++ utility classes.

Grid middleware support
The port-forwarding software installed in HPCx (dante) has been upgraded to a new version. This new version together with a number of local modifications has significantly improved the capabilities of this package. The port forwarded is now capable of successfully forwarding sockets created by the Globus toolkit libraries.

In particular it allows users to distribute applications across multiple sites using the MPICH-G2 package. This is a version of the MPI libraries built on top of the Globus toolkit. A number of users have expressed interest in using MPICH-G2 on HPCx. It is worth noting however that though it is now quite easy to compile programs to run across multiple sites only applications that have been explicitly optimised for the grid will be able to do the efficiently. This is because even dedicated high performance long distance networks like UKLight are significantly slower than the HPC interconnect available between the HPCx nodes. Message latency is also a major issue. We measured the MPICH-G2 latency between HPCx and a machine in Edinburgh at 4 milliseconds. Though this number is only a couple of times the absolute limit of latency imposed by the speed of light it is still several orders of magnitude greater than the MPI latency between HPCx nodes.

We have written a technical report explaining how users can use the port forwarder to run MPICH-G2 jobs on HPCx.

- **HPCxTR0606** "Grid Metacomputing support on HPCx", S. Booth

User administration software
The user administration software has been extended to allow registered users to view the results of the on-line surveys. The results page follows the same basic form as the original survey form. For each possible answer to a multiple choice question the results page will show

- The number of times that response was selected.
- The total number of responses to that question.
- The percentage of question responses for that response.

The responses are also colour coded by percentage. It is also possible to view and additional text comments submitted with the questionnaire.

We have also added an additional page to allow project managers to perform many more detailed analysis of the way their projects use time, displayed in different forms and broken down in various ways, including:

- User
- Budget
- Program
- Job size
- Discount

4.5 Operations and Systems *(Mr Mike Brown)*

4.5.1 Service development and planning

Simultaneous Multi-threading
Simultaneous Multi-threading (SMT) is a facility of the POWER5 processor, which allows it to operate as two virtual processors running two threads. Following preparatory work in the previous quarter, this has now been made available to users. We believe that we have achieved an implementation which is easy to use, and which can be enabled and disabled at will by each user. Changes were made to the job filter to recognise a simple flag in a user's job script to enable SMT; IBM provided a short program to enable the root SMT control program to be run in a user environment.

A full description can be found in the following Technical Note:

- **HPCxTR0604** *"An Investigation of Simultaneous Multithreading on HPCx"*, A. Gray, J. Hein, M. Plummer, A. Sunderland, L. Smith, A. Simpson, A. Trew

Phase 3 preparation
The plan for the move to Phase 3 is still on track. During a standard maintenance session in April, IBM carried out important preliminary work on the switch cabling.

UKLight
The appropriate interfaces on the login node and the routing information required have been set up to enable this to be used, and it was activated in June.

4.5.2 Incidents

There were 4 SEV1 incidents (contractual failures) this quarter, and 34 other incidents, mostly at SEV3.

Staffing

AV	*April*	*May*	*June*
DL	4.9	4.4	4.8
EPCC	6.9	8.7	8.0
Total	11.8	13.1	12.8

Systems	5.2	6.2	5.8

5 Summary of Performance Metrics

Metric	*TSL*	*FSL*	*April*	*May*	*June*
Technology serviceability	80%	99.2%	99.7%	99.3%	100.0%
Technology MTBF (hours)	200	300	732	732	∞
Number of AV FTEs	7.5	10	11.8	13.1	12.8
Number of training days per month	22.5/12	30/12	7/4	20/5	20/6
Non in-depth queries resolved within 3 days	85%	97%	98.6%	100.0%	99.1%
Number of A&M FTEs	3.75	5.75	5.2	6.2	5.8
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 Trubulence Consortium	Dr Gary Coleman
e05	1	Materials Chemistry using Terascaling Computing	Prof Richard Catlow
e06	1	UK Car-Parrinello Consortium	Prof Paul Madden
e07	2	Turbulent Plasma Transport in Tokamaks	Dr Colin M Roach
e08	2	Organic Solid State	Prof Sarah Price
e10	1	Reality Grid	Prof Peter Coveney
e11	1	Bond making and breaking at surfaces	Prof Sir David A King
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
e19	1	Edinburgh Soft Matter and Statistical Physics Group	Prof Michael E Cates
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
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
e27	1	SPICE	Prof Peter V Coveney
e29	1	Free-surface-piercing circular cylinders	Dr Eldad Avital
e30	1	Metal/Oxide Interfaces at the Atomic Level	Dr Nora de Leeuw
e31	1	Lateral Straining of Wall-Bounded Turbulence	Dr Gary N Coleman
e32	1	Rapid Prototyping of Usable Grid Middleware	Prof Peter V Coveney
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
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
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
e48§	1	Organised structure in turbulent flows	Prof Sergei Chernyshenko
e49§	1	Integrated Programme of Research in Aeronautical Engineering	Prof Michael Leschziner
e50§	1	Biological interface with materials	Prof John Harding
e51§	1	Super-computing data mining	Dr Mike Pettipher
e52§	1	Spacecraft force modelling	Dr M Ziebart
e53§	1	Large-scale communication networks	Prof J M Pitts
e54§	1	Free surface simulation of waves overtopping during storms	Dr D M Ingram

e55§	1	High-Reynolds-Number Near-Wall Flows	Prof Michael Leschziner
z09		HECToR Benchmarking	Dr Edward Smyth

§ Project transferred from CSAR, June 2006

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*
b02	1	Modelling enzyme catalysis	Dr Adrian J Mulholland
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 J-C Desplat
x03	IBM	Mr Derrick J Byford

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

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
e09	2	Molecular Properties and their Geometry	Prof Peter Taylor
e12	1	Parallel programs for the simulation of complex fluids	Dr Mark R Wilson
e13	1	TeraGyroid project	Dr Richard J Blake
x02		OHM Ltd	Mr Mark Westwood
n05	2	Non-linear Wave-particle Instabilities in Plasmas	Dr Mervyn Freeman