HECToR Annual Report 2011

01 January – 31 December 2011

Issue: 1.0

Contents

1 Introduction

This report covers the period from 1 Jan 2011 at 08:00 to 1 Jan 2012 at 08:00.

The next section of this report contains an Executive Summary for the year.

Section 3 summarises service availability, performance and utilisation statistics for the year. Section 3 also covers the Helpdesk statistics. Systems support is covered in Section 4, with the work of the Cray Centre of Excellence described in Section 5 and the Computational Science and Engineering (CSE) Support provided by NAG covered in Section 6.

The Appendices define some of the terminology and incident severity levels and list the current HECToR projects together with their overall utilisation profile to date.

This report and the additional SAFE reports are available to view online at http://www.hector.ac.uk/about-us/reports/annual/2011.php

2 Executive Summary

2011 was another busy period for the HECToR Service. The highlights of the service over the year included:

- Average utilisation on the XE6 in 2011 was just over 70%. With the expansion to Phase 3, there is unallocated capacity on the service looking forward into 2012/2013. Revised allocation models are being proposed which should help to address the projected shortfall in allocations.
- The Phase 2b system (XE6) formed the contractual service as from February 2011. Both the Phase 2a system (XT4) and X2 were decommissioned in May. Utilisation on both the XT4 and XE6 combined was 69%. Usage on the XT4 decreased steadily when the XE6 was introduced.
- The upgrade to the final phase of HECToR took place in late 2011. A key aspect of the Phase 3 upgrade was to minimise cost by building on existing hardware and re-using existing infrastructure. The system was upgraded to a 30 cabinet Interlagos based system. This equates to 90,112 cores, with a peak capacity of over 800 T/Flops.
- There were 10 technology-attributed service failures in 2011, compared to 45 in 2010. In addition to the technology failures, there were 3 site failures and two external failures. The overall MTBF was 586 hours compared to 195 hours in 2010. The XE6 has proven to be significantly more reliable than the XT4.
- The NERC Large Memory server (LMS) was installed in February. This is now used for pre and post processing activities.
- The HECToR GPU test system was installed in March. Over 7000 jobs have been run, and the system was 50% utilised.
- The transition to the second external shared lustre filesystem went ahead in June. All non-NERC data was migrated to the secondary filesystem.There is now over 1PB of online storage.

- A total of 5401 queries were handled in 2011 and the associated Helpdesk statistics for the year were excellent. For the first year since HECToR has been in service, no negative quality tokens were received.

- The CSE training programme continued to be popular with many courses being full to capacity. A number of new courses were introduced in 2011 to reflect changes to the service, and training materials revised in preparation for Phase 3. To date there have been over 1200 attendees on training courses, with excellent feedback received.

- By the end of 2011, 54 years of effort had been allocated to improve the applications running on HECToR. The CSE Performance Working Group, which monitors the performance of the CSE service, met three times during 2011 and evaluated 18 DCSE projects which had recently completed. The average scores awarded were 4.6/5 for the extent to which they had achieved their objectives, and 4.3/5 for the impact that they had had. Improvements to codes are always fed back to the code owners so that the benefits of this work can be felt throughout the whole HPC community.

- The Cray Centre of Excellence (CoE) was heavily involved in preparing users for the arrival of the Phase 3 Cray XE6 Interlagos system. This involved both improving scaling of applications to take advantage of the greater number of cores, and assistance in the OpenMP hybridisation of certain applications to take advantage of the increase in cores on a node.

- An initial meeting with users to review a proposal for a HECToR tertiary storage system was held in July. Key HECToR user groups contributed their initial feedback on what users would want from such a storage system. The Invitation-To-Tender was issued in December with a target for installation by 31st March 2012.

3 Quantitative Metrics

3.1 Reliability

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

	1Q11	*2Q11*	*3Q11*	*4Q11*	*2011*
Incidents	50	92	42	41	225
Failures	5	5	1	4	15

The incidents above are primarily related to single node failures. Details on both the service failures and single node fails in 2011 can be found in Section 5.

3.1.1 Performance Statistics

- MTBF = (732)/(number of failures in a month)

  Quarterly MTBF = (3x732)/(number of failures in a quarter)

  Annual MTBF = (12*732)/(number of failures in a year)

Attribution	*Metric*	*1Q11*	*2Q11*	*3Q11*	*4Q11*	*2011*
Technology	Failures	5	3	0	2	10
	MTBF	439	732	∞	1098	732
Service Provision	Failures	0	1	0	2	3
	MTBF	∞	2196	∞	1098	1757
External/Other	Failures	0	1	1	0	2
	MTBF	∞	2196	2196	∞	4392
Overall	**Failures**	**5**	**5**	**1**	**4**	**15**
	MTBF	**439**	**439**	**2196**	**549**	**586**

In contrast to 2010, the system was much more reliable. This can be attributed to the upgrade to the XE6 which had much improved resilience over the XT4. In 2011 there were 10 technology failures, as opposed to 45 in 2010.

3.2 HECToR Utilisation

3.2.1 Phase 2b/Phase 3 Utilisation

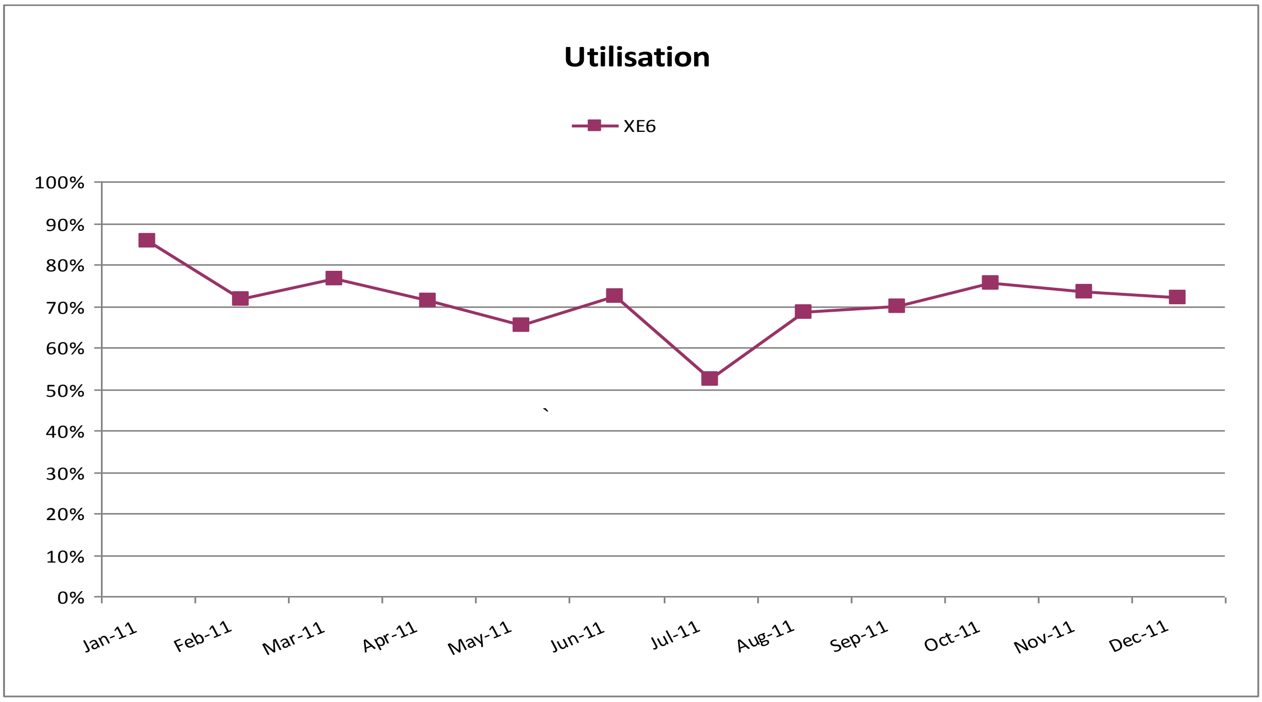

Overall utilisation of the XE6 in 2011 was 71%.

The system comprised a number of different phases over the year. In January the Phase 2a (XT4) formed the contractual service. This was superseded by the Phase 2b system (XE6) in February. The Phase 2a system was decommissioned in May. Finally, the service entered Phase 3 in December, with the introduction of the new Interlagos architecture and increased capacity on the service. Overall utilisation of the service on all phases in 2011 was 69%, compared to 73% in 2010.

The usage spilt across the various phases can be seen below.

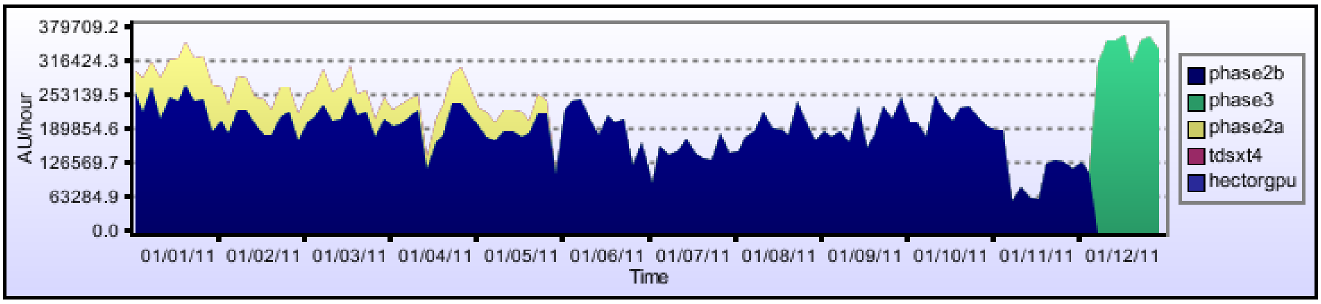

RawAUs by Machine in 2011

3.2.2 Phase 2a Utilisation

The Phase 2a (Cray XT4) system was decommissioned in May 2011. Utilisation for the 5 months averaged 52% winding down from over 70% to under 40% by the end.

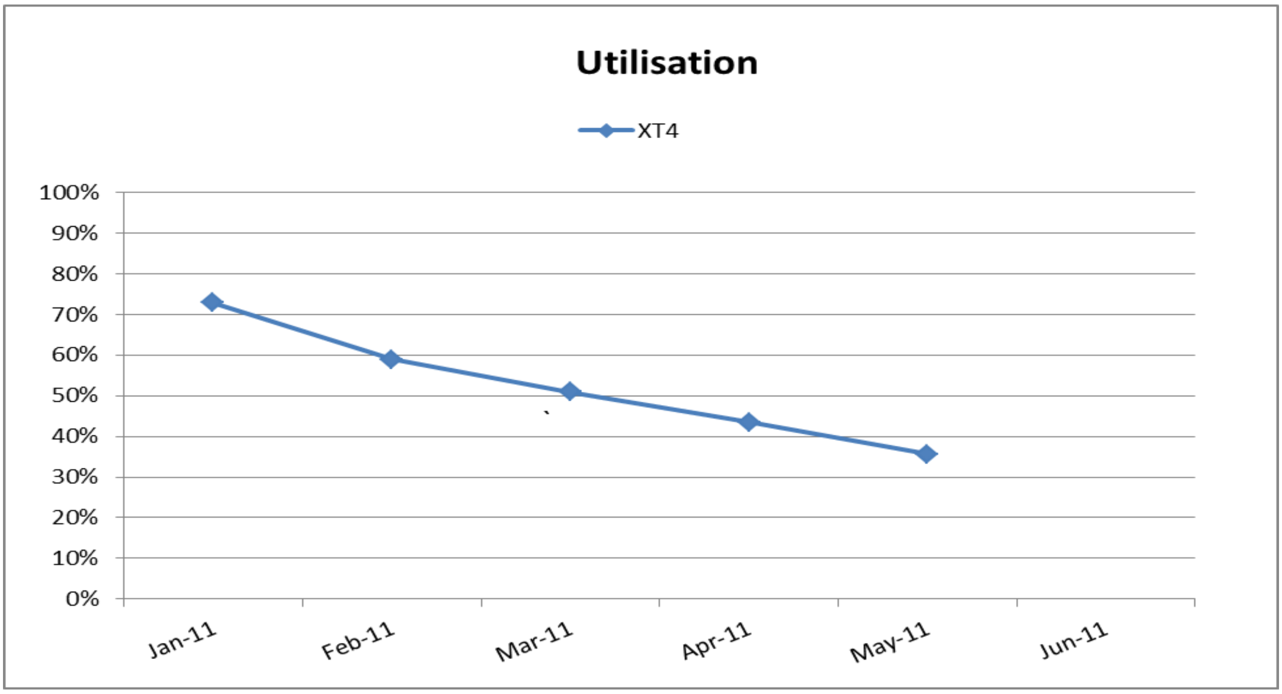

3.2.3 X2 Utilisation

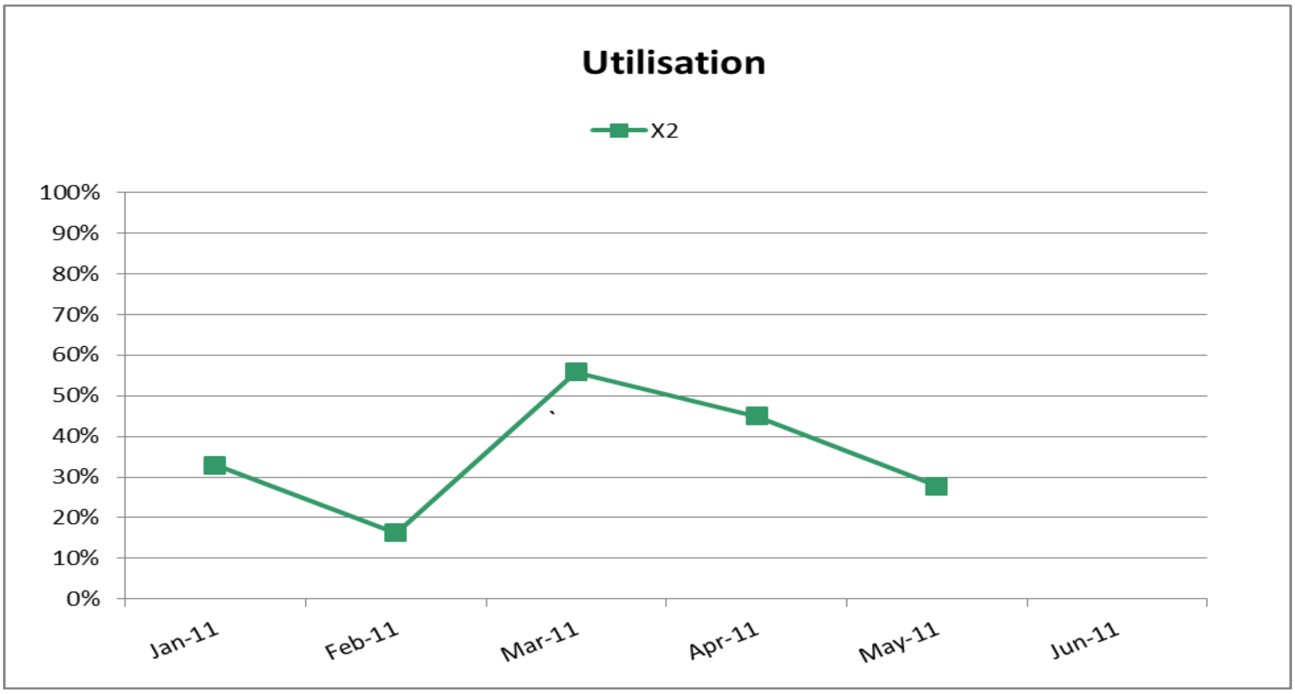

The X2 was decomissioned in May 2011. Utilisation for the 5 months was 37%.

3.2.4 HECToR Utilisation by Project Class

There are five main project classes on HECToR:

Class1a – Full Peer Review
Class 1b – Direct Access (or RAP)
Class 2a – Pump Priming
Class 2b – DCSE
Service – Support/Directors Time

The primary usage on HECToR is as expected from full peer reviewed 'Class 1a' projects.

Usage from Direct Access projects increased from just over 6% in 2010, to 7.7% in 2011. Whilst direct access projects should be for a set period only – typically 6 months, is it still common for requests to be received by the Helpdesk to extend them for a couple of months. In circumstances where service downtime has impacted the execution of the projects, these extensions have been approved by EPSRC.

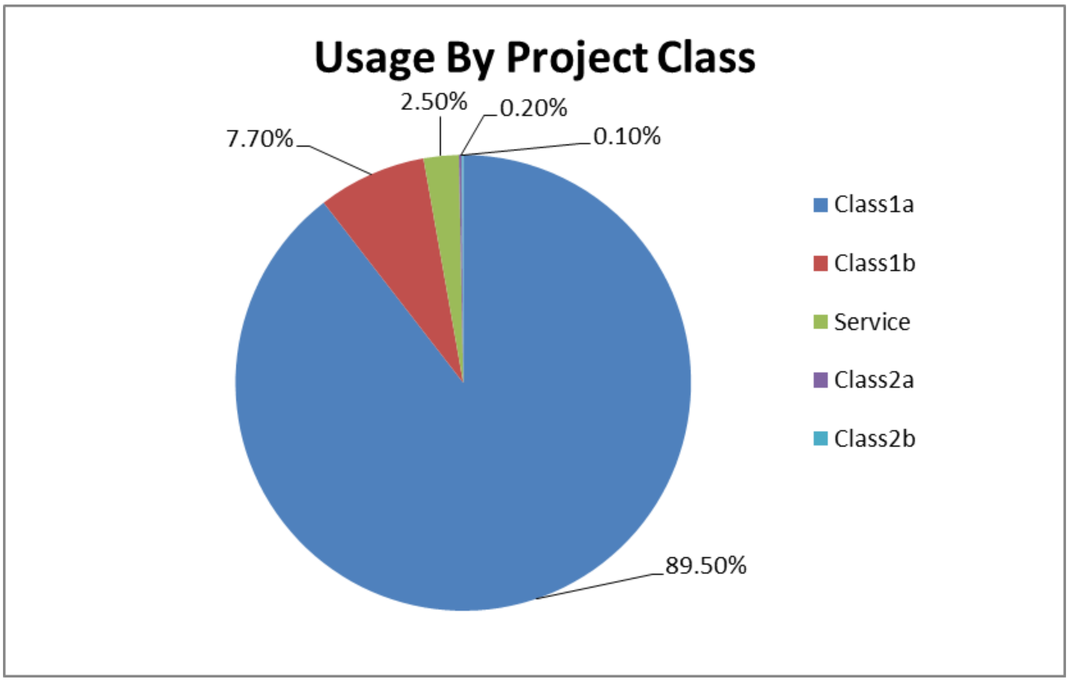

3.2.5 HECToR Utilisation by Consortium

As below, the main utilisation in 2011 came from the major EPSRC and NERC consortia as one would expect.

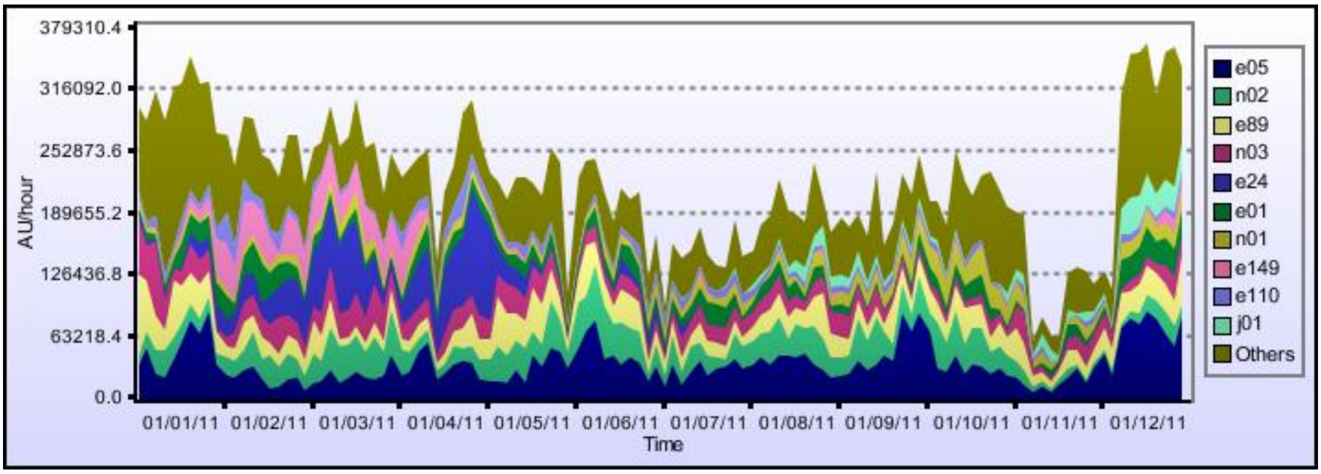

A total of 2,800,200.7 kAUs were available during this period. The utilisation across all phases was a follows:

Project	Raw kAUs	Number of Jobs	Raw %age	Utilisation
y01	0.5	9	0.0%	0.0%
y02	12270.9	1626	0.6%	0.4%
y03	384.5	2638	0.0%	0.0%
y04	0.0	1	0.0%	0.0%
y05	0.0	4	0.0%	0.0%
y06	0.1	832	0.0%	0.0%
y07	3.2	534	0.0%	0.0%
z01	563.4	2699	0.0%	0.0%
z02	0.2	9	0.0%	0.0%
z03	8765.8	7940	0.5%	0.3%
Internal Total	**21988.7**	**16292**	**1.1%**	**0.8%**
c01	14847.4	6665	0.8%	0.5%
e01	112152.4	7523	5.8%	4.0%
e05	330452.3	55687	17.1%	11.8%
e10	3936.3	831	0.2%	0.1%
e19	0.0	1	0.0%	0.0%
e24	127460.5	13163	6.6%	4.6%
e35	0.0	1	0.0%	0.0%
e42	282.2	54	0.0%	0.0%
e63	7629.1	131	0.4%	0.3%
e68	8389.6	2500	0.4%	0.3%
e70	499.3	48	0.0%	0.0%
e71	10379.3	1466	0.5%	0.4%
e72	0.0	1	0.0%	0.0%
e76	1556.1	581	0.1%	0.1%
e81	0.0	1	0.0%	0.0%
e82	6511.3	406	0.3%	0.2%
e84	742.3	38	0.0%	0.0%

Project	Raw kAUs	Number of Jobs	Raw %age	Utilisation
e85	25810.5	525	1.3%	0.9%
e89	199211.7	30445	10.3%	7.1%
e92	5198.9	958	0.3%	0.2%
e102	0.2	13	0.0%	0.0%
e104	5314.8	2101	0.3%	0.2%
e107	16056.9	717	0.8%	0.6%
e108	19094.8	1351	1.0%	0.7%
e110	63452.5	4228	3.3%	2.3%
e117	8915.4	9771	0.5%	0.3%
e120	22.7	21	0.0%	0.0%
e121	94.0	5	0.0%	0.0%
e122	34067.2	6495	1.8%	1.2%
e124	24745.3	648	1.3%	0.9%
e125	17444.3	709	0.9%	0.6%
e126	5086.2	827	0.3%	0.2%
e127	1596.5	77	0.1%	0.1%
e128	4780.3	335	0.3%	0.2%
e129	380.1	82	0.0%	0.0%
e130	4830.8	253	0.3%	0.2%
e131	158.0	15	0.0%	0.0%
e134	148.1	136	0.0%	0.0%
e135	3.4	21	0.0%	0.0%
e136	275.4	124	0.0%	0.0%
e137	0.0	1	0.0%	0.0%
e139	241.0	521	0.0%	0.0%
e141	6264.3	768	0.3%	0.2%
e144	0.3	20	0.0%	0.0%
e145	233.4	129	0.0%	0.0%
e146	222.9	78	0.0%	0.0%
e147	39.4	10	0.0%	0.0%
e148	435.5	409	0.0%	0.0%
e149	69297.4	531	3.6%	2.5%
e152	23570.6	224	1.2%	0.8%
e155	202.3	35	0.0%	0.0%
e156	72.6	45	0.0%	0.0%
e157	99.8	116	0.0%	0.0%
e158	1182.0	774	0.1%	0.0%
e159	15.4	338	0.0%	0.0%
e160	1.2	199	0.0%	0.0%
e167	31172.4	932	1.6%	1.1%
e170	2142.7	247	0.1%	0.1%
e171	1255.3	224	0.1%	0.0%
e173	265.7	597	0.0%	0.0%
e174	269.0	43	0.0%	0.0%
e175	20322.1	1160	1.1%	0.7%
e177	389.7	142	0.0%	0.0%
e178	0.9	16	0.0%	0.0%
e179	884.6	229	0.1%	0.0%
e182	0.0	7	0.0%	0.0%
e183	0.0	1	0.0%	0.0%
e184	1742.6	1271	0.1%	0.1%
e185	7527.3	283	0.4%	0.3%
e186	26289.3	3821	1.4%	0.9%

Project	Raw kAUs	Number of Jobs	Raw %age	Utilisation
e187	19.5	53	0.0%	0.0%
e188	2355.7	850	0.1%	0.1%
e189	5997.8	580	0.3%	0.2%
e190	5684.1	443	0.3%	0.2%
e191	4569.8	1060	0.2%	0.2%
e192	4555.3	366	0.2%	0.2%
e193	24184.6	1294	1.3%	0.9%
e194	5204.0	142	0.3%	0.2%
e195	6565.1	889	0.3%	0.2%
e196	4735.2	62	0.3%	0.2%
e198	66.3	69	0.0%	0.0%
e199	1561.5	676	0.1%	0.1%
e201	156.0	74	0.0%	0.0%
e203	446.6	379	0.0%	0.0%
e205	177.6	177	0.0%	0.0%
e208	250.9	275	0.0%	0.0%
e210	227.0	204	0.0%	0.0%
e211	1.4	41	0.0%	0.0%
e212	2180.8	79	0.1%	0.1%
e213	0.0	1	0.0%	0.0%
e215	6535.0	525	0.3%	0.2%
e216	4411.1	2543	0.2%	0.2%
e217	2042.2	1378	0.1%	0.1%
e218	126.3	253	0.0%	0.0%
e219	1589.3	90	0.1%	0.1%
e222	555.9	228	0.0%	0.0%
e224	76.4	34	0.0%	0.0%
e225	106.3	37	0.0%	0.0%
e227	59.5	149	0.0%	0.0%
e228	192.3	32	0.0%	0.0%
e230	200.9	158	0.0%	0.0%
e231	0.0	12	0.0%	0.0%
e234	2.9	10	0.0%	0.0%
e238	0.0	2	0.0%	0.0%
e239	0.0	5	0.0%	0.0%
j01	39453.6	2256	2.1%	1.4%
EPSRC Total	**1379957.0**	**177551**	**71.5%**	**49.3%**
n01	91468.1	20537	4.7%	3.3%
n02	205707.7	109535	10.7%	7.4%
n03	155528.7	28027	8.1%	5.6%
n04	25705.9	25510	1.3%	0.9%
NERC Total	**478410.4**	**183609**	**24.8%**	**17.1%**
b10	567.4	564	0.0%	0.0%
b12	0.0	1	0.0%	0.0%
b13	132.1	205	0.0%	0.0%
b14	36.2	21	0.0%	0.0%
b100	0.0	2	0.0%	0.0%
BBSRC Total	**735.6**	**793**	**0.0%**	**0.0%**
p01	133.7	475	0.0%	0.0%
STFC Total	**133.7**	**475**	**0.0%**	**0.0%**
x01	12878.3	10799	0.7%	0.5%
x05	2.8	55	0.0%	0.0%
x06	0.1	68	0.0%	0.0%
External Total	**12881.2**	**10922**	**0.7%**	**0.5%**
d03	494.7	356	0.0%	0.0%

Project	Raw kAUs	Number of Jobs	Raw %age	Utilisation
d04	368.2	5215	0.0%	0.0%
d11	1322.8	274	0.1%	0.1%
d15	26.7	226	0.0%	0.0%
d16	191.0	1109	0.0%	0.0%
d18	10.8	25	0.0%	0.0%
d19	1498.2	56	0.1%	0.1%
d22	3.1	21	0.0%	0.0%
d23	204.8	685	0.0%	0.0%
d24	0.5	16	0.0%	0.0%
d25	7304.6	2241	0.4%	0.3%
d26	51.0	3247	0.0%	0.0%
d27	38.6	179	0.0%	0.0%
d28	15622.6	5850	0.8%	0.6%
d29	0.5	20	0.0%	0.0%
d30	111.4	66	0.0%	0.0%
d32	176.3	861	0.0%	0.0%
d33	11.2	32	0.0%	0.0%
d34	14.9	417	0.0%	0.0%
d35	0.1	5	0.0%	0.0%
d37	951.3	1095	0.1%	0.0%
d39	26.8	220	0.0%	0.0%
Director’s Time Total	**28430.0**	**22216**	**1.5%**	**1.0%**
pr1u0705	1.9	1	0.0%	0.0%
pr1u0706	6592.1	121	0.3%	0.2%
PRACE Total	**6594.0**	**122**	**0.3%**	**0.2%**
Total	**1929130.7**	**411980**	**100.0%**	**68.9%**

3.2.6 HECToR Utilisation by Application Area

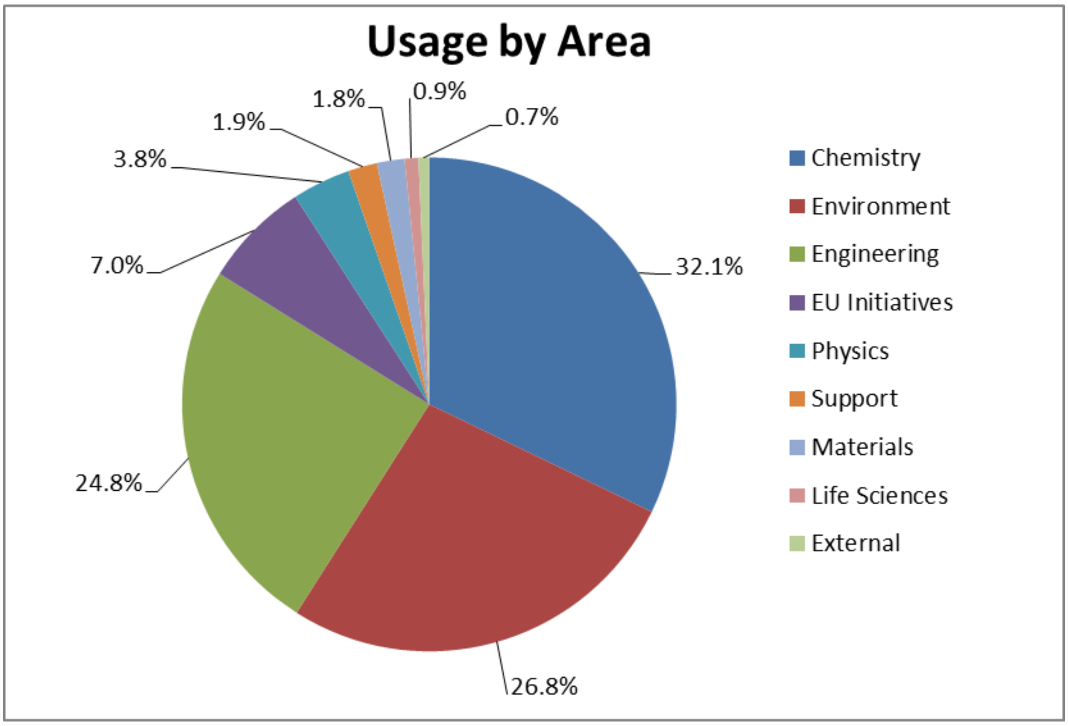

3.2.7 HECToR Code Usage Statistics

The chart below shows the ten codes accounting for the most utilisation on HECToR.

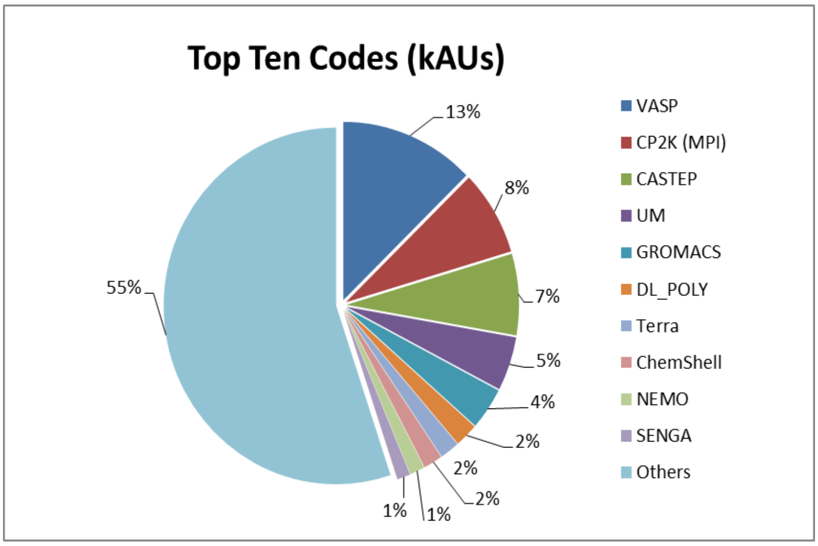

Code	% kAUs	% Jobs
VASP	12.65%	17.61%
CP2K (MPI)	7.73%	6.35%
CASTEP	7.37%	5.68%
UM	4.90%	7.83%
GROMACS	3.82%	4.31%
DL_POLY	2.17%	3.18%
Terra	1.84%	0.81%
ChemShell	1.78%	0.94%
NEMO	1.41%	0.59%
SENGA	1.24%	0.05%
ONETEP	1.18%	1.06%
NAMD	0.97%	3.90%
CPMD	0.73%	1.40%
CP2K (Hybrid)	0.67%	0.38%
CASINO	0.57%	0.04%
LAMMPS	0.57%	0.57%
Quantum Espresso	0.54%	1.80%
Shelf	0.36%	0.41%
Fluidity	0.14%	0.15%
Amber (PMEMD)	0.10%	0.27%
HELIUM	0.07%	0.01%
SIESTA	0.05%	0.18%
GAMESS-US	0.04%	0.04%
Others	49.10%	42.44%

3.2.8 HECToR Projected Utilisation

The chart below shows a projection of system use. The chart is a best guess of the projected future use of the system based on the per project profiles. If the project has over or under used their time allocation relative to the projected profile then the profile is re-scaled to take account of this. For dates in the past the actual use is displayed.

The steep drop shown in the chart from April is partially due to the NERC allocation method. NERC allocate budget to their various consortia on an annual basis from April-April and the 2012-2013 allocation has not yet been confirmed. Despite this it can be seen that with the increase in capacity there is still a shortfall in allocation.

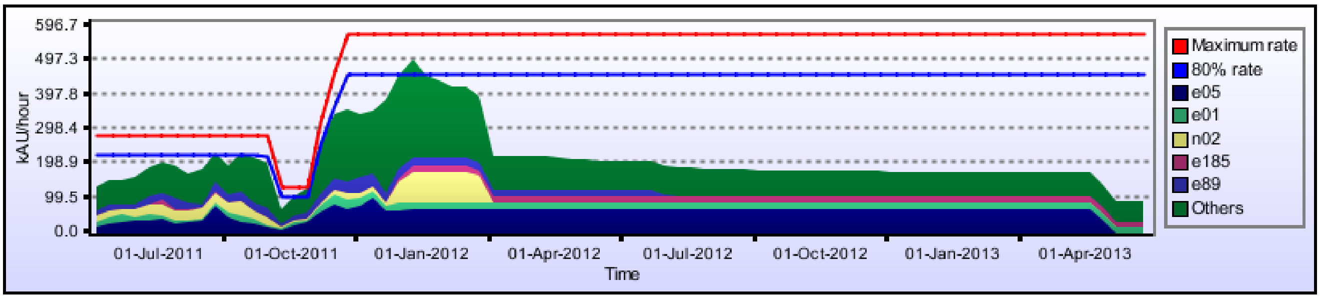

An action had been taken previously by EPSRC to review and renew allocations on the commencement of a new Phase. The initial proposal will be reviewed in late February.

3.3 HECToR GPU

The HECToR GPU test system came online in March. The initial allocation was shared 50% EPSRC and 50% Edinburgh. In total there have been just 4 EPSRC/NERC projects which have been setup on the GPU as part of the Class1b access mechanism. In addition a number of Edinburgh based projects have trialled codes.

From June to January the system was 50% utilised. In total, over 7000 jobs have been run on the system. The purpose of the test system is to test using GPUs as opposed to running major production runs. As such we expect to see a high number of short test jobs which may not necessarily use many GPU hours.

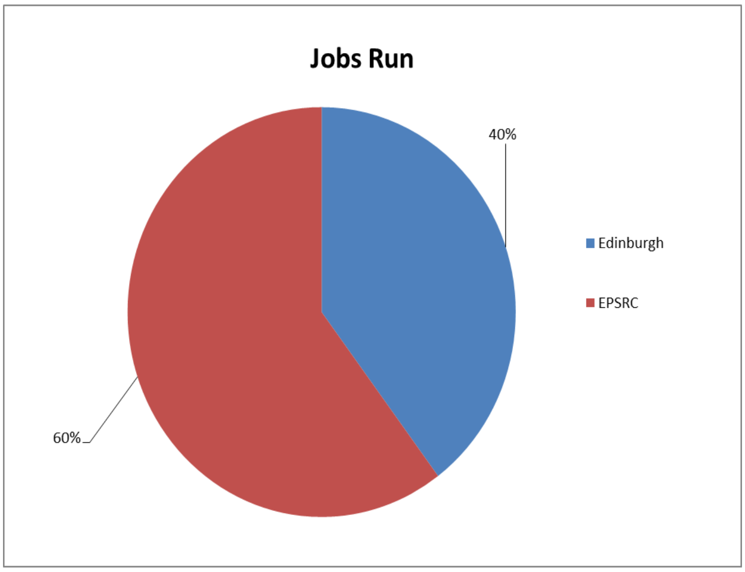

Whilst the split in the volume of jobs has been fairly even, the GPU usage has been dominated by one EPSRC project which is using a code already ported for the GPU.

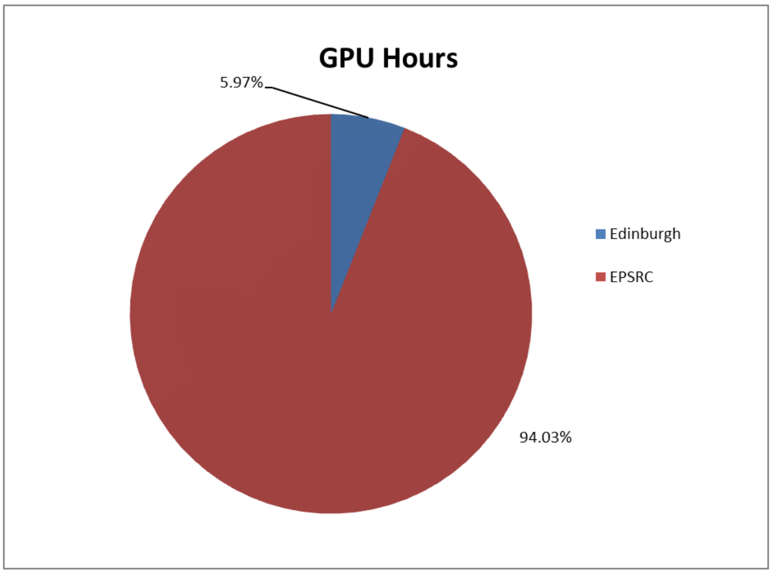

More users should be encouraged to use the GPU in 2012.

3.4 Helpdesk

A total of 5401 queries with a specified service metric were completed in this period. There were also 103 queries with no metric completed in the same period.

3.4.1 Helpdesk Targets

Metric	Pass	Total	Fraction	Target
All queries finished in 1 day	4611	4661	98.9%	97.0%
Admin queries finished in 1 day	4274	4316	99.0%	97.0%
Queries assigned in 30 min	5339	5366	99.5%	97.0%
Technical assessments in 10 days	70	75	93.3%	97.0%

3.4.2 Queries by Service Metric

Service Metric	Queries	Percentage
Automatic	3206	58.2%
Admin	1110	20.2%
In-depth	665	12.1%
Technical	345	6.3%
No Metric	103	1.9%
Technical assessment class-1b	31	0.6%
Technical assessment class-2a	20	0.4%
Technical assessment class-1a	20	0.4%
Technical assessment class-2b	4	0.1%

3.4.3 Queries by Category

Query Category	Queries	Percentage
Set group quotas	771	14.3%
Set user quotas	757	14.0%
New Password	651	12.1%
New User	603	11.2%
3rd Party Software	367	6.8%
Disk, tapes, resources	286	5.3%
Access to HECToR	238	4.4%
Batch system and queues	195	3.6%
User Behaviour	193	3.6%
None	192	3.6%
Compilers and system software	188	3.5%
Login, passwords and ssh	112	2.1%
Node Failure	110	2.0%
User programs	107	2.0%
Add to group	81	1.5%
Other	75	1.4%
Join Project	74	1.4%
New Group	73	1.4%
Make Reservation	65	1.2%

Query Category	Queries	Percentage
Remove account	48	0.9%
Update account	47	0.9%
SAFE	33	0.6%
Courses	32	0.6%
Archive	22	0.4%
Delete from group	17	0.3%
Performance and scaling	15	0.3%
Static website	10	0.2%
Network	10	0.2%
Create certificate	10	0.2%
Delete from project	9	0.2%
Porting	6	0.1%
Grid	4	0.1%

3.4.4 Queries by Handler Category

Handlers	Total	Automatic	Admin	In-depth	Technical	TA	No Metric	Percentage
OSG	3448	3205	115	26	90		12	62.6%
USL	1429	1	985	166	193		85	26.0%
Cray	198		9	127	57		5	3.6%
CSE	429		2	346	5	75	1	7.8%

3.4.5 Helpdesk General

3.4.5.1 esFS

During June and July, the helpdesk played a major communication role during the migration of data to the secondary external lustre filesystem esFS. Data for all non-NERC projects was migrated to the new filesystem. This work was split into 9 phases. The phased transition of data for the ten projects with large volumes of data required co-ordination and communication with the individual user groups.

3.4.5.2 Phase 3

The HECToR website was updated in November for the arrival of the Phase 3 system.

3.4.5.3 HECToR iPhone Application

In March 2011 the HECToR User iPhone application was launched. This is now also available on the Android platform. This enables users to check the system status, and queue status while on the move. Users can also be notified automatically when their jobs complete. Across the two platforms there were over 250 downloads of the app in 2011.

3.4.5.4 Code Porting Activities

The contractual applications as agreed with EPCC were ported to the Phase 3 system in November.

3.4.6 User Quality Tokens

No negative tokens were received in 2011. In total 74 positive tokens were set, compared to 45 positive tokens in 2010. A summary per project is below:

Project	***Positive Tokens***
e05	27
n01	16
x01	15
e139	5
c01	5
d23	4
e10	2
Total	**74**

3.5 Performance Metrics

Metric	*TSL*	*FSL*	*January*	*February*	*March*	*April*	*May*	*June*	*July*	*August*	*September*	*October*	*November*	*December*	*Annual Average*
Technology reliability (%)	85.0%	98.5%	99.3	98.6	100.0	97.1	99.1	100.0	100.0	100.0	100.0	100.0	94.5	99.9	99.1
Technology MTBF (hours)	100	126.4	244.0	366.0	∞	366.0	732.0	∞	∞	∞	∞	∞	732.0	732.0	878.4
Technology Throughput, hours/year	7000	8367	8625	8475	8648	8326	8267	8557	8647	8621	8784	8615	7102	8588	8483
Capability jobs completion rate	70%	90%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%
Non in-depth queries resolved within 1 day (%)	85%	97%	100.0%	99.4%	98.7%	99.3%	99.7%	97.3%	97.3%	98.5%	98.7%	98.7%	99.0%	98.8%	98.7%
Number of SP FTEs	7.3	8.0	8.2	8.3	9.4	8.1	10.1	10.0	8.5	8.5	9.1	8.8	8.5	7.4	8.7
SP Serviceability (%)	80.0%	99.0%	100.0%	100.0%	100.0%	100.0%	100.0%	99.7%	100.0%	100.0%	100.0%	100.0%	97.7%	100.0%	99.8%

Colour coding:

Exceeds FSL	(green)
Between TSL and FSL	(amber)
Below TSL	(red)

4 Systems Hardware

4.1 HECToR Technology Changes

2011 was another very busy period for the service, with a number of key technology changes taking place. Service disruption was kept to a minimum whenever possible.

4.1.1 Phase 2b Acceptance Testing

The year began with the acceptance testing of the XE6 system in January. The tests were completed successfully and the Phase 2b system became the contractual service as of February.

4.1.2 NERC LMS

The NERC Large Memory Server (LMS) was installed and tested in February. A subset of NERC users was given access to the server to perform post-processing activities.

4.1.3 HECToR GPU

The HECToR GPGPU test-bed completed acceptance tests in March with initial user groups online as of April.

4.1.4 X2/XT4 end of Life

The Phase 2a systems (both XT4 and X2) were decommissioned at the end of May 2011.

4.1.5 esFS Secondary Filesystem

Work commenced in June to migrate all non-NERC data to a secondary external Lustre filesystem. The migration plan was formulated to cause as little disruption as possible. This work was completed as planned on 10 July.

4.1.6 CLE Upgrade

In preparation for the move to Interlagos, the operating system was upgraded to CLE4.0 in November.

4.1.7 Phase 3 Upgrade

The HECToR service was upgraded to an Interlagos system in late 2011. A key aspect of the upgrade was to minimise cost by building on existing hardware and re-using existing infrastructure. The system was upgraded to a 30 cabinet Interlagos based system. This equated to 90,112 cores, with a peak capacity of over 800 T/Flops.

The upgrade work was planned in three phases to minimise disruption:

- Phase One – Reduce capacity of XE6 to 10 cabinets
- Phase Two - Return 20 cabinet Interlagos based system
- Phase Three – Return 30 cabinet Interlagos based system

The system capacity was reduced as planned in early November. It was identified that additional testing was required before the 20 cabinet system could be brought into user service. This caused some delay to the original plan, with the upgraded system being brought online on 7th December. The final 10 cabinets were then added in January, with acceptance testing and availability trials completed successfully thereafter.

4.2 Severity-1 Incidents

4.2.1 Technology Failures

Cray supplied technology has been responsible for 10 severity-1 incidents in 2011.

This is a breakdown of the failure categories:

- 3 cabinet power/control/cooling component failures
- 3 esFS server/controller hardware or firmware failures
- 1 lustre software error
- 1 CLE HSN software error
- 2 instances of a late return to service following scheduled planned maintenance sessions

Reliability has improved significantly during 2011. The principle reason for the improvement has been due to the benefits of the Gemini Interconnect fabric, which includes robust resiliency features and is implemented in the Phase 2b/Phase 3 (Cray XE6) system.
All three of the cabinet related incidents occurred during January when the Phase 2a (Cray XT4) system was still the contracted service, utilising the previous generation “Seastar” Interconnect fabric.

The three esFS related hardware or firmware errors were caused by one disk controller, one lustre OSS server and one firmware error in a disk controller. A second esFS filesystem was introduced this year which helps minimize the risk of such incidents from affecting all users.

The software incident caused by a lustre error was resolved by an upgrade to lustre 1.8.4 in November.

The CLE High Speed Network related software incident was fixed by the installation of a patch in April.

The two incidents of a late return to service following scheduled maintenance sessions are unfortunate but considerable effort is made to thoroughly plan and prepare for these sessions to prevent such incidents from occurring. Improvements in processes and procedures continue to be made to minimize these risks.

4.2.2 Service Provision Failures

There were three failures attributed to the service provision in 2011.There was one occurrence of a safety system tripping. In November there were two cooling related incidents, which occurred during the expansion of the Phase 3 system. The necessary changes to the configuration were applied to prevent a repeat failure.

4.2.3 External Failures

There were two other failures in 2011. There was one external power failure, and one network failure.

4.3 Single Node Failures

The number of single node fails causes no concern. The peak of node fails in May was attributed to two related software bugs which were addressed in a subsequent maintenance update.

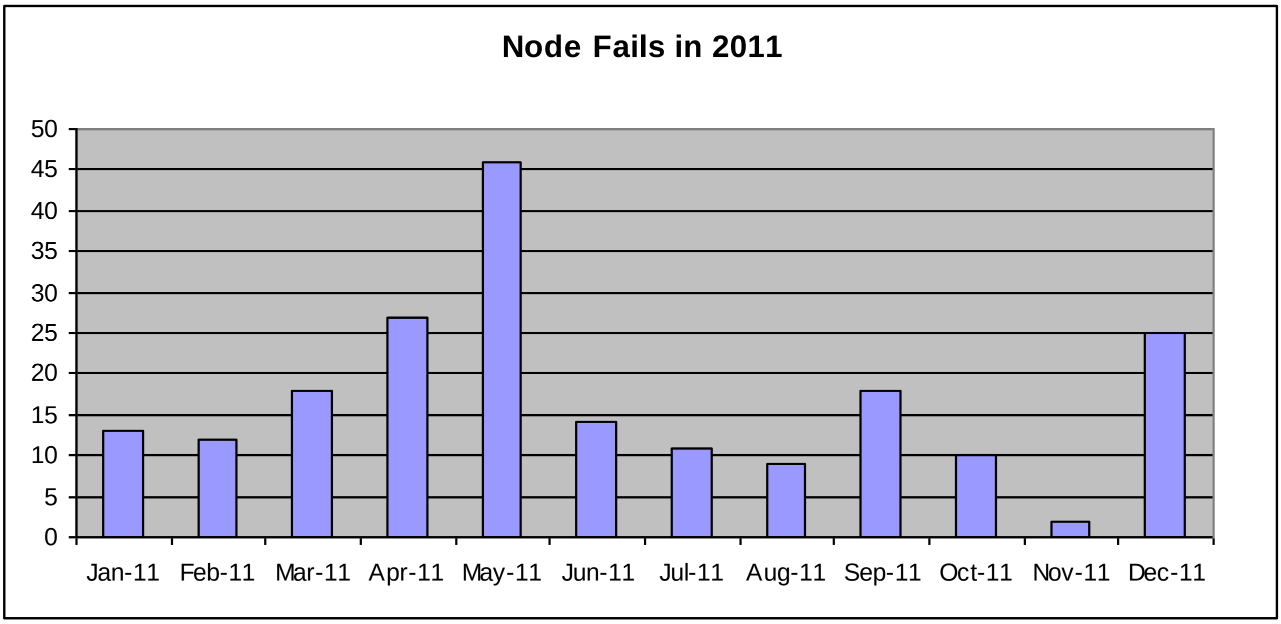

5 Cray Centre of Excellence

5.1 Executive Summary

The year 2011 saw the Cray Centre of Excellence for HECToR continue to deliver targeted high-end application support to a wide range of user groups. The Centre was heavily involved, along with the wider Cray applications team, in preparing users for the arrival of the Phase 3 Cray XE6 Interlagos system. This involved both improving scaling of applications to take advantage of the greater number of cores, and assistance in the OpenMP hybridisation of certain applications to take advantage of the increase in cores on a node.

The Centre has engaged in user training and workshops, promoting the HECToR service throughout the UK, Europe and the rest of the world. Members have contributed to the distributed CSE process and Performance Working Groups; provided guidance and support for postgraduate students participating in High Performance Computing training; and contributed to the preparation of the UK's HPC user community and application base for the arrival of HECToR Phase 3 system. This was in addition to the continued involvement of the Centre's core staff and members of Cray's Application and Development teams with HECToR activities where specialist input was essential.

5.2 Applications Support

- The Cray CoE for HECToR closely collaborated with NERC's National Centre for Atmospheric Sciences Computation Modelling Support team in preparing the Unified Model for the HECToR upgrade to Phase 3. This was an extensive effort to both port the UM suite from Magny Cours to Interlagos combined with switch of default compiler from Pathscale to the Cray Compiling Environment (CCE). A large cross section of model configurations were required to pass fundamental functionality tests when compiled with CCE, whilst the most important configurations were subject to an extended test suite. A key part of this test suite was ensuring "bit reproducible" results from UM configurations that required it. This work stretched over several months, involving close collaboration between the Cray CoE for HECToR, Cray R&D personnel from the compiler and performance tools teams, and the NCAS CMS team in Reading. In total, over 20 UM model configurations and versions were ported successfully to CCE, which has now become the default compiler of choice for the Unified Model user community on HECToR. This work has acted as a pioneer for other UM projects on other Cray systems to use CCE, including the on-going UPSCALE PRACE project on the Cray XE6, HERMIT, at HLRS.

- On-going support for OpenFOAM on HECToR was also provided by the Cray CoE for HECToR. Early in the year OpenFOAM ports for GEMINI were provided, and more recently OpenFOAM ports for Interlagos have also been provided. In addition work was done to better understand the scalability of OpenFOAM on the Cray XE6. Figure 4 below shows the scalability, measured as parallel speedup, of OpenFOAM when simulating lid-driven cavity flow up to 8k cores. Running the same test case on the

Cray XT5 the scalability dropped off after 2k cores, so this demonstrates nicely the improved strong scaling possible with GEMINI.

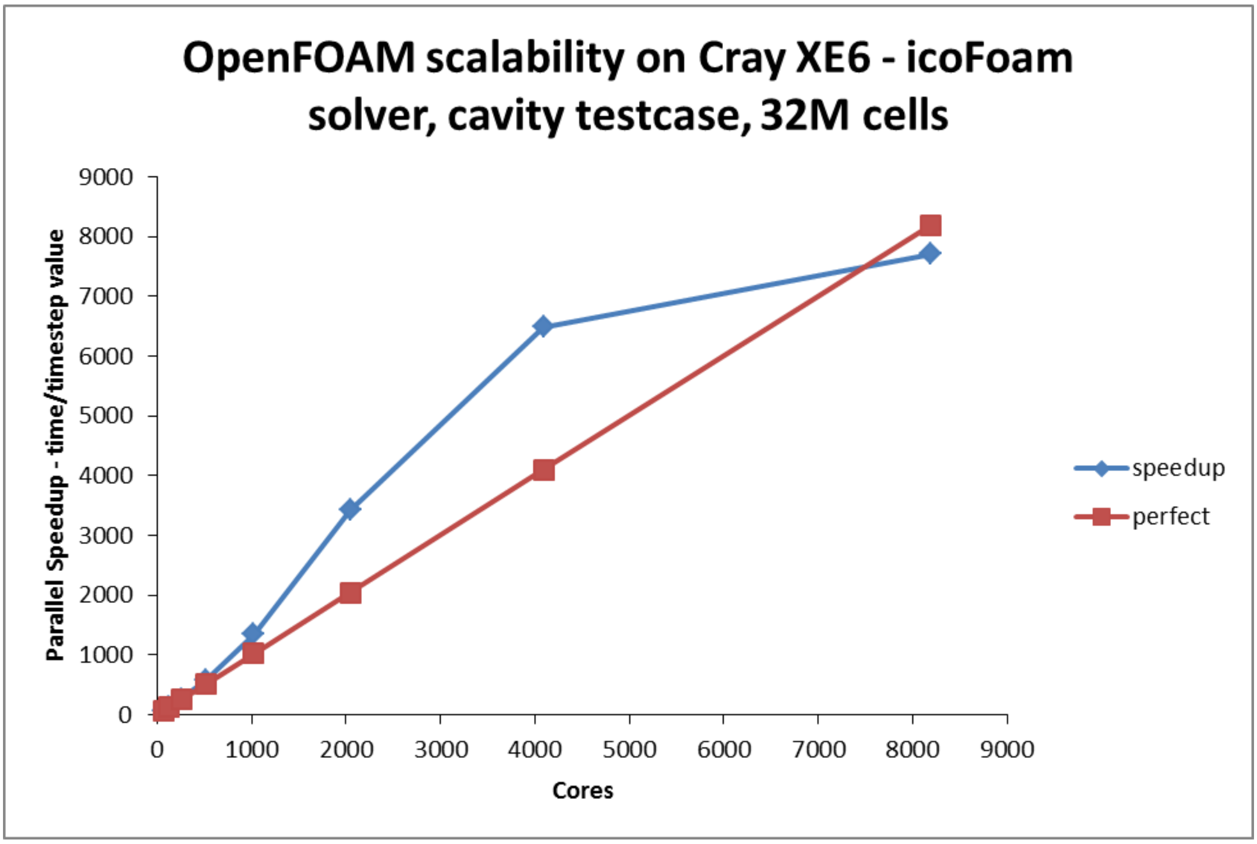

- HiPSTAR is a CFD simulation of subsonic jet engines developed by Dr Richard Sandberg at the University of Southampton. HiPSTAR's original parallel decomposition was restricted to two dimensions using MPI which limited the total number of processor cores any problem could use. The increased number of shared memory cores in each node of the Cray XE6, from 4 to 24 allowed further decomposition along the third dimension using OpenMP. The hybrid code has been tested on the HECToR's Cray XE6, comparing the pure MPI version with the hybrid. Each version ran on the same total number of cores, either with an individual MPI rank per core or with one MPI rank decomposed over 6 threads (the internal structure of the Opteron Magny-Cours processor favours 4x6 threads compared to 1x24 threads). Runs ranged from 384 (16 nodes) cores to 24 576 (1024 nodes – the largest queue size on HECToR) all with the same problem size (strong scaling). The first result was that the pure MPI task did not complete when running on 24 576 cores. The hybrid version ran to completion demonstrating that decomposing along the extra dimension allows the application to scale to larger core counts. Figure 2 shows that the hybrid version has better scaling and continues to scale further than the original pure MPI.

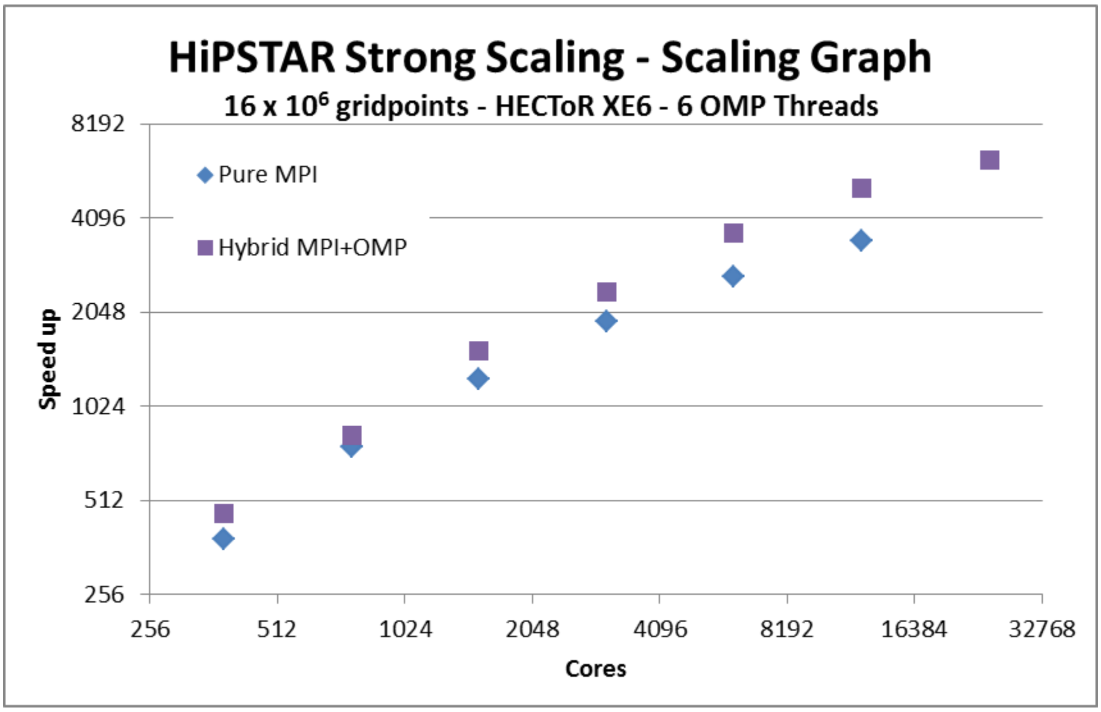

Figure 1 HiPSTAR Scaling Graph

Figure 3 shows the total computational cost on all processors for each time step. It shows that for all core counts the hybrid code runs more quickly on the same total number of cores resulting in a faster time to solution. The improvements range between 10% - 50% gains, with, importantly, the greatest gains at the largest core counts.

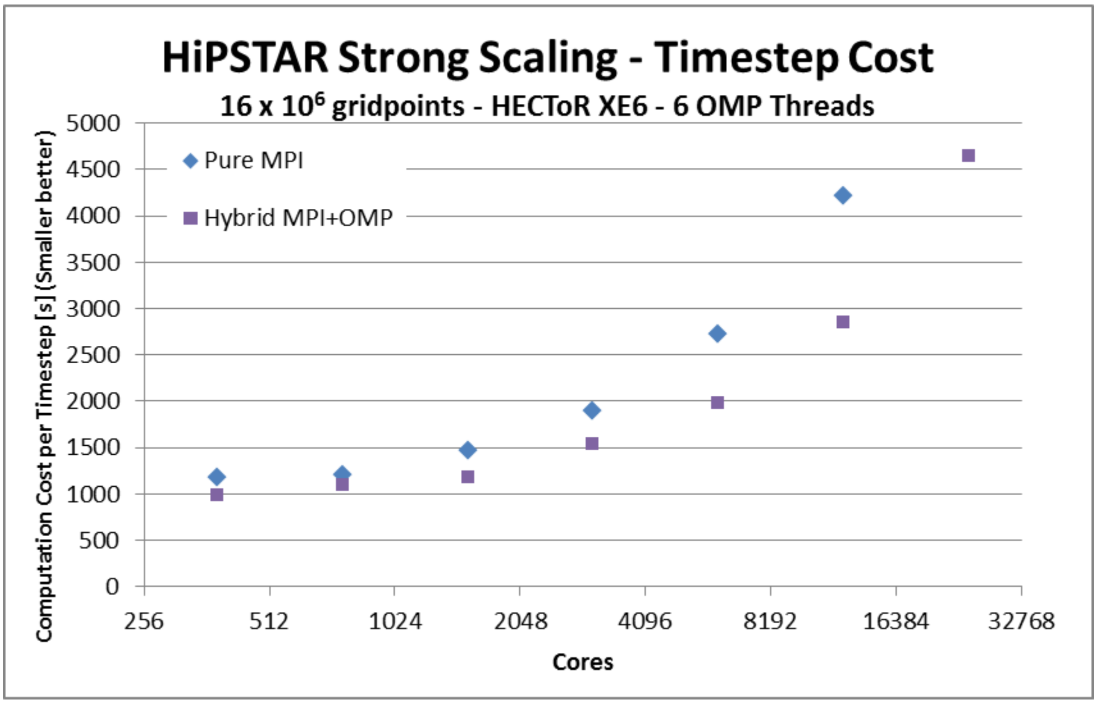

Figure 2 HiPSTAR Timestep Cost Graph

5.3 Research Projects

- The upgrade to the Gemini interconnect has enabled installation of Cray's "Cluster Compatibility Mode (CCM)" which is a key technology for users of industrial tools. Applications supplied by Independent Software Vendors which have been targeted at generic Linux clusters can now be used on Cray XE6s' like HECToR. The CoE, combined with the Cray system's team have assisted the installation and configuration of ANSYS, a suite of CFD tools which is highly popular in industry, using the CCM mode on the TDS. This has allowed an industrial user of ANSYS to evaluate the software on the Cray XE6 with the aim of large-scale engineering simulations on the production system.

- During 2011 a joint research project between Cray Exascale Research Initiative Europe, the Cray CoE for HECToR, Cray Storage personnel, and HPCX to look at Lustre monitoring tools and approaches was undertaken. The aim was to investigate existing tools for monitoring Lustre IO traffic from a server side perspective and find out which approach/toolset would work best in the HECToR environment. This work was carried out on the HECToR TDS system with the aim of future implementation of the favoured approach on the production system. After some initial work looking at various software packages for doing this, it was decided to focus on the Lustre Monitoring Tools (LMT) as these are in use at other Cray sites, and come with much of the functionality required here at HECToR. The end result of this work was very successful study showing both what could be collected with LMT on HECToR, but also what was required in terms of software and hardware to do this collection. A report detailing the software developed to do this collection, installation and usage was also produced. Further work was also done to find out what could be collected on a per-user basis in order to track IO traffic at a fine-grained user level.

5.4 Workshops and Training Events

The Cray CoE for HECToR conducted a number of training events in 2011:

- Members of the Cray CoE for HECToR provided a training workshop at NAG's Oxford campus in February focusing on the upgraded Gemini interconnect and recent upgrades to the Cray software environment.

- At the DEISA/PRACE Spring School in Edinburgh, 29-31 March, the Cray CoE for HECToR gave a one-day tutorial on using the Cray XE6 and optimization for the Cray XE6. This was followed by an afternoon hands-on session working with users and their applications.

- The Cray CoE for HECToR, in conjunction with the Cray CoE at HLRS for the HERMIT system, gave a three-day user workshop, 2-4 February, on using the Cray XE6 at HLRS.

- The Cray CoE for HECToR gave a one-day workshop on OpenFOAM porting and optimization to users at HLRS on 3 February.
- Visited several user groups throughout the year to meet with a larger group of users, give updates on the HECToR service, and give targeted application support where necessary. This included visits to the UK Turbulence Consortium and Materials Chemistry Consortium meetings, along with several visits to the NCAS researchers in Reading.

5.5 Seminars

The Cray CoE for HECToR has given a number of external seminars of interest to the UK HPC community:

- ICOMEX workshop: The Cray CoE for HECToR was invited to represent HECToR and Cray at the inaugural meeting of the ICOMEX project at DWD, Offenbach in November. ICOMEX is an international research project to investigate future use of Icosahedral grids in weather modeling to which the Cray CoE for HECToR was able to provide a technical insight.
- HPCN Workshop: The Cray CoE for HECToR and the Exascale Research Initiative Europe gave a presentation titled “Hybrid Multicore Supercomputing and GPU Acceleration with Next-Generation Cray Systems” at the HPCN Workshop in Braunschweig, 25-26th May. This presentation covered the Cray GPU programming environment, and more generally programming for multicore and heterogeneous compute nodes.
- SimGPU 2011: The Cray CoE for HECToR and Exascale Research Initiative Europe gave a presentation titled “Hybrid Multicore Supercomputing and GPU Acceleration with Next-Generation Cray Systems” at the HPCN Workshop in Braunschweig, 31 May-1 June. This presentation covered the Cray GPU programming environment, and more generally programming for multicore and heterogeneous compute nodes.

6 The HECToR Computational Science and Engineering (CSE) Support Service

6.1 Overview of the CSE Service

The Computational Science and Engineering (CSE) service exists to help the user community to make the best use of the HECToR hardware by providing training, web-based resources, and assistance with porting, optimisation and tuning of software. The service is provided by the Numerical Algorithms Group Ltd (NAG), a not-for-profit Company with offices in Oxford and Manchester, and 40 years' experience developing mathematical and statistical software. The *Core Team*, made up entirely of NAG staff, responds to in-depth software problems reported by users via the HECToR helpdesk, processes Technical Assessments related to applications for access, runs a range of training courses and maintains a range of good practice guides and reference material as part of the service website, and undertakes various outreach activities. The *Distributed Team*, made up of a mixture of NAG staff and staff employed via a contract with a third party, provides dedicated resources for projects (DCSE projects) to enhance specific applications or support groups of users for periods of between three months and a year.

6.2 Highlights of 2011

2011 was a busy year for the service and the CSE Team. The XE6 system came into service at the end of 2010 and in May the X2 and XT4 systems were finally decommissioned, while a GPU test-bed system was introduced in March. The CSE team has helped users with all of these changes, both directly through the helpdesk where we dealt with more than 500 in-depth queries, through providing updated reference material on the website, and via appropriate training courses.

The CSE training programme continues to be popular with many courses being full to capacity. This year we introduced a number of new courses to reflect changes to the service, and we have revised all our training material in preparation for Phase 3. To date we have filled over 1200 places on our training courses, and received excellent feedback from attendees.

The DCSE programme has gone from strength to strength. By the end of 2011, fifty-four years of effort had been allocated to improve the applications running on HECToR. The CSE Performance Working Group, which monitors the performance of the CSE service, met three times during 2011 and evaluated 18 DCSE projects which had recently completed. The average scores awarded were 4.6/5 for the extent to which they had achieved their objectives, and 4.3/5 for the impact that they had had. Improvements to codes are always fed back to the code owners so that the benefits of this work can be felt throughout the whole HPC community.

6.3 The CSE Helpdesk

The Core CSE team handles queries from users forwarded by the service helpdesk, carries out technical assessments of applications for HECToR time, undertakes various outreach activities and runs the training courses.

The queries received by the CSE team vary from straightforward requests for advice to requests for assistance in porting, tuning etc. Some queries are resolved straight away while others develop into small projects lasting weeks or even months. The team resolves most queries but if, after investigation, they are found to be connected to system issues, then they will be re-assigned to the Service Provider (UoE HPCx) or to Cray. The following table shows how many queries were opened and closed each month.

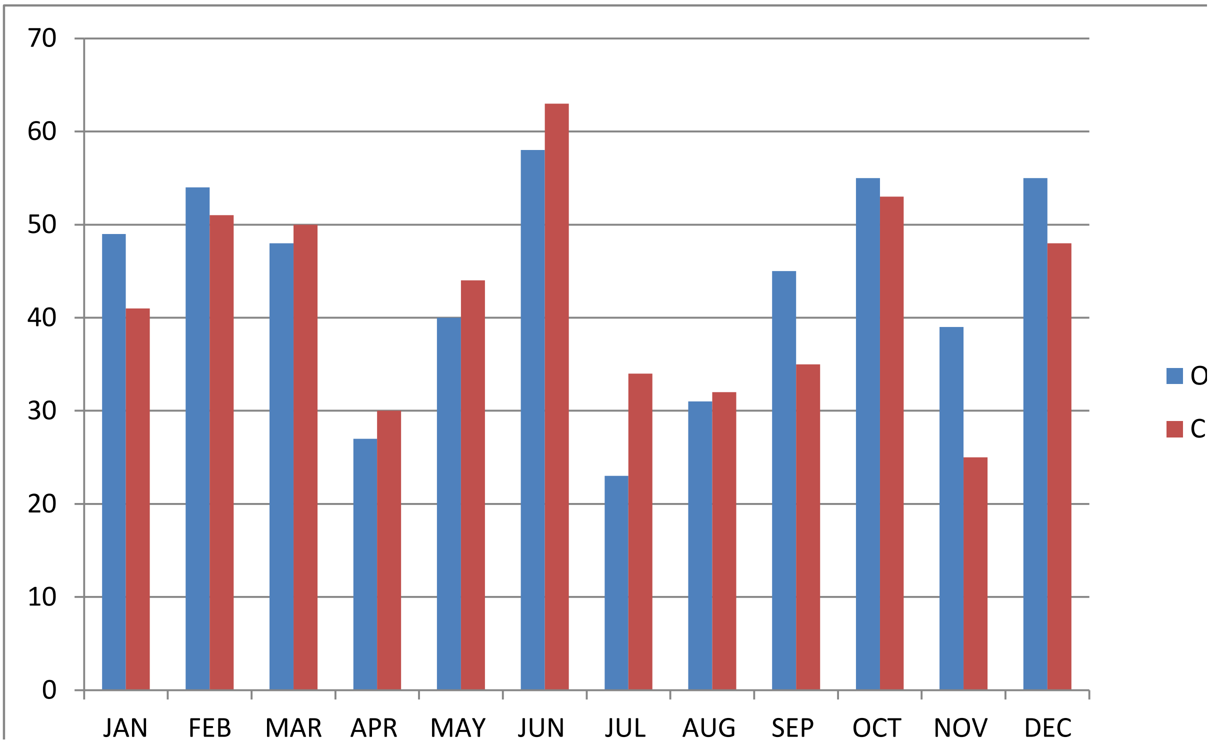

In cases where a technical query (as opposed to a request for a Technical Assessment) is resolved by the team (strictly speaking where the query is closed within the SAFE system by the CSE team), the user is invited to fill in a questionnaire giving feedback about his or her experience and satisfaction with the outcome. This year, 124 have been returned. On the rare occasions that a negative response is received in feedback the CSE team will attempt to understand the reasons behind the response and, if necessary, change or improve their procedures. Most feedback is, however, extremely positive.
The responses to the questionnaire are summarised as follows:

Has the problem raised in your query been resolved by the information provided by the helpdesk?

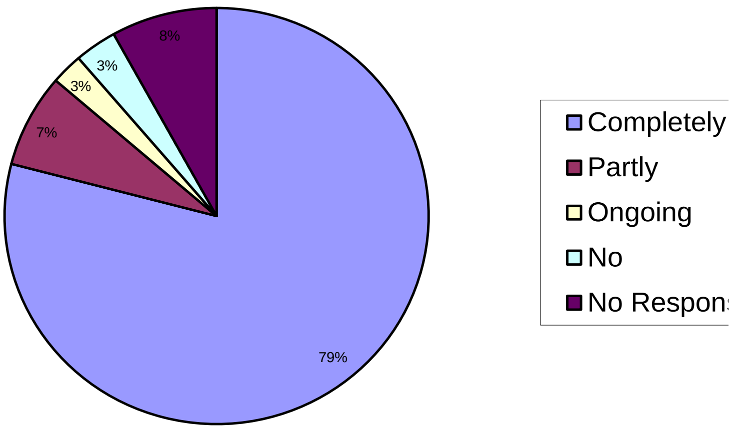

Were all communications replied to promptly and the text cle and understandable?

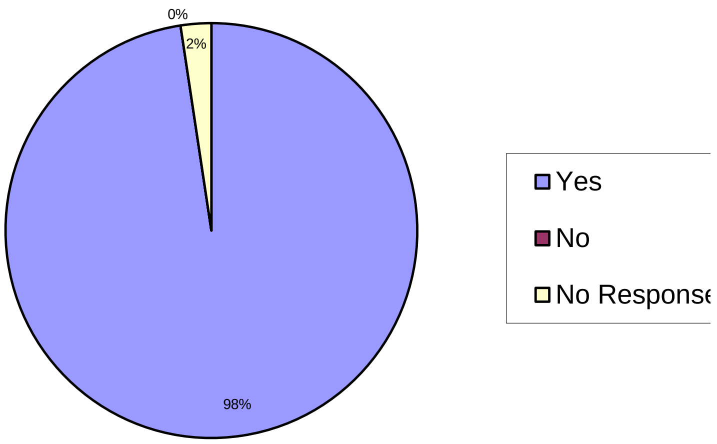

Were technical issues explained in sufficient detail?

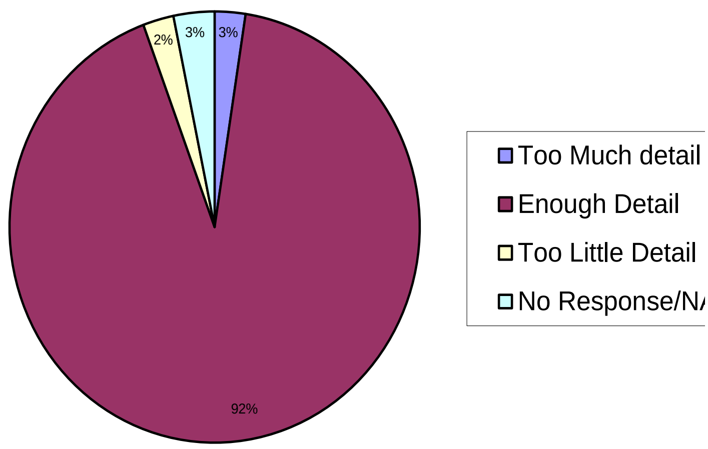

Did you use the HECToR CSE documentation to try and find a solution before submitting a query to the helpdesk?

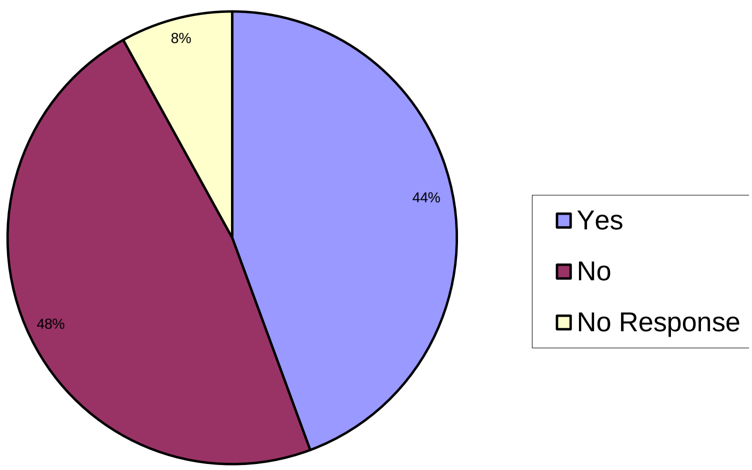

6.4 Training

During this period NAG has offered a range of training courses, most of which were not HECToR-specific. To support the GPU test-bed we introduced courses in CUDA and OpenCL; to keep-up with changes in Fortran we ran courses on co-arrays and Object-Oriented programming; and with the decommissioning of the XT4 we ran courses to help those users who had not already done so to migrate to the XE6. Most courses have been run by NAG staff, although we have delivered courses in collaboration with the Universities of Warwick, Bristol, Oxford and Edinburgh, and with SANDIA Laboratories.
Take-up for training this year has been extremely good, with many of the courses being delivered at users' institutions. We have provided training at Imperial College, Bath, Warwick, King's College London, UCL, Southampton and Reading, as well as at the NAG offices in Oxford and Manchester. A total of 71 days of training have been offered and delivered to 419 attendees.

The complete list of courses offered in 2011 is as follows. The courses range in duration from one day to a week, and complementary courses are often scheduled together for convenience. Many of these courses were run several times in different locations. In addition we delivered a module on Advanced Computational Methods within the Doctoral Training Centre for Complex Systems at the University of Southampton, and ran two workshops during a NERC summer school.

- Parallel Programming with MPI
- OpenMP
- Introduction to HECToR
- Multicore
- Fortran 95
- Debugging, Profiling and Optimising
- Core Algorithms for High Performance Scientific Computing
- Transitioning to the Cray XE6
- Co-array Fortran
- Object-oriented programming in Fortran 2003
- An Introduction to CUDA Programming
- An Introduction to OpenCL Programming
- SPRINT: Parallel Computing with R on HECToR

6.5 Other Core CSE Activities

The CSE team makes regular meetings to users, both individually (for example to provide assistance preparing an application for HECToR time or discuss training needs), and collectively (for example attending meetings of the UK Turbulence Consortium). We attend appropriate conferences and workshops, and maintain links with other major international HPC facilities.

6.6 Distributed CSE

The panel met twice during 2011, and considered a total of twenty one applications requesting 170 months of effort. Of those, fourteen have been recommended to receive a total of 94 months of effort.

Eighteen projects ended during 2011. At the end of December fifteen projects were running with five more scheduled to start in 2012. Since the start of the service the DCSE panel has allocated just over 54 person-years of effort. In October NAG organised a workshop for all DCSE participants in Manchester, slides from which are available on the HECToR website.

6.7 Distributed CSE Projects Completed during 2011

6.7.1 Parallel algorithms for the materials modelling code CRYSTAL

CRYSTAL is an *ab initio* electronic structure and materials properties code, which is widely used on HECToR by groups at Bath, Cambridge, Imperial, Kent, Southampton, UCL and Warwick. This DCSE project was allocated 24 months support at the January 2008 call. Professor Nic Harrison of the Department of Chemistry at Imperial College London supervised the work, performed by Stanko Tomic from STFC between February 2009 and February 2011.

The aim of the project was to reduce the memory bottleneck by improved memory optimisation and code modularisation. The replicated data structure of the Hartree-Fock and density matrix were removed and replaced with their irreducible representations. Speed-ups in multiples of problem size/(2 times the symmetry of the system) can now be achieved. Reducing the data block size from the default value of 96 to 64 or 32 showed more than a 10% speed-up in the diagonalisation part of the code, and a 20% speed up in the back and similarity transform for up to 3,584 cores on Phase 2b. This work has been introduced within the main code base.

6.7.2 Parallelisation of CABARET

The CABARET (Compact Accurate Boundary Adjusting high REsolution Technique) code may be used to solve the compressible Navier-Stokes equations, as required by Large Eddy simulation (LES) models of aircraft noise. This DCSE project was allocated 12 months effort at the December 2008 round. Phil Ridley of NAG worked on the project between March 2009 and January 2011.

The aim of this project was to develop a scalable version of the compressible CABARET code. A parallel partitioning method for the structured partitioning of the unstructured CABARET grid was developed, along with the facility to use ParMetis or PT-Scotch for future geometries. An MPI parallel version of CABARET was also implemented using non-blocking MPI to pass data between cell faces and sides. The new code demonstrated a parallel efficiency of 72% on 250 cores of Phase 2a. A hybrid (OpenMP/MPI) parallel version of CABARET was then developed for Phase 2b. For a 3D backward-facing step test case of 51.2M cells, a parallel efficiency of 80% was achieved on 1,000 cores of Phase 2b.

Simulations can now be performed on grids at least 512 times larger than was previously possible. For smaller grid sizes, time averaged high Reynolds number (Re) simulations can now be performed in hours rather than days. The project PI, Dr Sergey Karabasov of the Department of Engineering at the University of Cambridge, holds around 50M HECToR AUs, of which over half will be used for CABARET simulations. Since February 2011 around 7.5M AUs have been used for CABARET on Phase 2b.

6.7.3 Direct Numerical Simulations (DNS) of turbulent fluid flows

Incompact3D is used to conduct cutting-edge turbulence studies within the group of Professor John C. Vassilicos, of the Department of Aeronautics at Imperial College London. This DCSE project was allocated 16.7 months of support at the December 2008 round.

Professor Vassilicos supervised the work, performed by Ning Li of NAG between March 2009 and February 2011.

The aim of the project was to improve the scalability of the MPI parallel Incompact3D on HECToR. Incompact3D's communication routines were re-written by updating the original domain decomposition from a 1D (slab) to a 2D (pencil) domain decomposition for the FFTs and improving the global transpositions with a shared memory approach to the MPI_Alltoallv(s). A shared memory Fortran library package for 3D Cartesian FFT solvers, 2DECOMP&FFT, was also developed. This is now available to all HECToR users, for scientific applications that may require upgraded distributed FFT calculations, e.g. combustion, ocean modelling and compressible CFD. It has already been used in several other CSE and DCSE projects.

A typical high resolution turbulence, mixing and flow control simulation involving a cubic mesh of size 4096 is now able to utilise 16,384 cores efficiently, so that productive scientific studies can now be done within 3.75 days rather than 25 days (wall-clock time). All improvements have been incorporated into the main Incompact3D release. Between October 2009 and June 2011 Incompact3D used 13.1M AUs on Phase 2a and, between December 2010 and October 2011, 77.1M AUs on Phase 2b.

6.7.4 Implementation of established algorithms to extend HELIUM

HELIUM generates accurate full-dimensionality solutions of the non-relativistic time-dependent Schrödinger equation for two-electron atoms or ions exposed to intense laser radiation. It is also one of the HECToR benchmark codes and is used within the group of project PI Professor Ken Taylor of the Department of Applied Mathematics & Theoretical Physics at QUB. This DCSE project was allocated 44 months effort at the December 2008 round. Michael Lysaght and Laura Moore of QUB worked on the project from April 2009 to March 2011 (Dr Lysaght took over when Dr Moore was seconded to another HELIUM project). Jonathan Parker, also of QUB, worked on the project from October 2009 to March 2011. Ed Smyth of NAG also worked on the project between May 2010 and March 2011. This project has extended the capabilities of HELIUM to enable it to fully exploit HECToR's capability. A generalized version of HELIUM for crossed and circularly polarized laser fields was developed and tested. HELIUM now has the unique ability to integrate the helium-laser Schrödinger Equation in its full generality (essentially 6-spatial dimensions plus time integration). This feature has already been used to search for failures in reduced dimensionality models of atom-laser physics. A parallel post-processing code, CYLINDRICAL, was also developed: this gives up to a 20 times speed-up for transforming the final-state wavefunction from configuration to momentum space. Additionally, the new *ab initio* computer code, RMT (R-matrix Incorporating Time), was also developed. RMT may be used for calculating the electron dynamics occurring within many-electron atoms when they are exposed to intense short-pulse light.

Professor Taylor's group has used almost 50M AUs on HECToR with the improved HELIUM and RMT codes. The RMT code is also the primary research tool of a collaboration between groups at The Open University, UCL, and QUB where another 130.5M AU will be used during Phase 3.

6.7.5 Hybrid time-dependent density functional theory in CASTEP Part 2

This was a continuation from a previously successful DCSE project, with the overall aim of implementing Time-Dependent Density Functional Theory (TDDFT) in the density functional theory code CASTEP and, therefore, increasing its usage on HECToR. This project was allocated 12 months of support at the December 2008 round. Dr Keith Refson of STFC supervised the work, performed by Dominik Jochym, also of STFC, from April 2010 to March 2011.

An XC response kernel was implemented for the hybrid functionals, enabling CASTEP to compute excited state energies of molecular systems using B3LPY, PBE0 or other hybrid functionals. The parallel TDDFT code in CASTEP was developed further to include the computation of forces in an excited state, and geometry optimisation of molecules in an excited state. Finally, Born Oppenheimer molecular dynamics was implemented to enable calculations of a molecular system in an excited state.

A scalable implementation of TDDFT in CASTEP will now give the UK electronic structure community an opportunity to address cutting-edge scientific problems in areas such as inorganic and organic photovoltaic materials, catalytic reactions at surfaces, light-emitting polymer materials for optical displays, and femtosecond laser chemistry. Including the use of hybrid functionals promises to address some of the limitations that have previously hindered the application of TDDFT to extended systems. Between October 2009 and June 2011, the centrally installed version of CASTEP alone consumed 64.2M AUs on Phase2a and 385.8M AUs on Phase2b between December 2010 and October 2011. This is not taking into account the usage by people who compile their own version of CASTEP.

6.7.6 Dynamic load balancing and rigid body dynamics for DL_POLY_3

DL_POLY is a general purpose package for classical molecular dynamics simulations, developed by Dr I.T. Todorov and Professor W. Smith of STFC. This DCSE project was allocated 18 months of effort at the March 2009 round. Dr Todorov supervised the work which was performed by Laurence Ellison, who was also based at STFC, from October 2009 to March 2011.

The DL_POLY_2 or DL_POLY_CLASSIC classical molecular dynamics package uses a full force field and molecular description. However, scaling is limited to tens of processes with up to 30,000 atoms. The successor, DL_POLY_3 was scalable for up to 512 processes with 1M atoms, but had no rigid body (RB) functionality. The aim of this project was to bring the RB functionality to DL_POLY_3 and also to implement dynamic load-balancing for further scalability. The RB functionality was implemented successfully and the resulting code is DL_POLY_4. A prototype load balancing algorithm was tested as a stand-alone code, but this was not introduced into the main DL_POLY_4 code base due to the unforeseen complexity of the task. However the results of performance tests carried out on Phase 2b were encouraging, with good speed ups observed for a maximum of 2744 cores with around 5M atoms. It is hoped that the load-balancing algorithm can be integrated within DL_POLY_4 in the future, as this would improve performance on many-core architectures. Between October 2009 and June 2011, the centrally installed version of DL_POLY_4 consumed 12.6M AUs on Phase 2a and 28.9M AUs on Phase 2b between December 2010 and October 2011. It is impossible to say how much of this use involves RB dynamics, but the known users of the RB work are at UCL, Bath, Cambridge, Warwick and Leicester.

6.7.7 Adding molecular dynamics functionality to CASINO

CASINO is used for performing quantum Monte Carlo (QMC) electronic structure calculations for finite and periodic systems. This DCSE project was allocated 18 months of support at the March 2009 round. The work was supervised by Professor Dario Alfè of the Department of Earth Sciences at UCL. Norbert Nemec, also based at UCL, worked on the project from January 2010 to August 2010. Dr Mike Towler of the University of Cambridge, who is also the main developer of CASINO, continued the work from October 2010 to July 2011.

The objective of this project was to couple the CASINO code to the molecular dynamics (MD) code PWSCF. The result was a scalable model for use on HECToR which will enable new scientific simulations to be carried out within the fields of both MD and QMC. Between October 2009 and June 2011, the centrally installed version of CASINO consumed 3.5M AUs on Phase 2a and 10.7M AUs on Phase 2b between December 2010 and October 2011.

6.7.8 Metal Conquest: efficient simulation of metals on petaflop computers

Conquest is a linear scaling code designed to perform Density Functional Theory (DFT) calculations on very large systems of atoms. It is currently used on HECToR by the group of Dr David Bowler of the Department of Physics and Astronomy at UCL, who supervised this DCSE project. The project was allocated 12 months effort at the September 2009 round. Lianheng Tong, also of UCL, performed the work between February 2010 and January 2011. The aim of the project was to develop Conquest so that it could be used efficiently for both insulators and metals. The Hamiltonian diagonalisation process was optimised by using smaller block sizes in ScaLAPACK which also improved load balancing, giving a 10% speed up. K-point parallelism was implemented using a BLACS process grid and a 21% speed up was demonstrated on a 32 atom 13x13x13 Al unit cell. Kerker and wave dependent metric preconditioning techniques as used in VASP were also implemented. For a bulk Al unit cell calculation with 32 atoms, Kerker preconditioning allows improved convergence. These developments will be committed to the central Conquest repository.

Around 5M AUs per year are consumed by Conquest at UCL, with further usage at Imperial College London. Additionally, a current joint BBSRC project with Dr Bowler has an allocation of 5.5M AUs, for the direct use of this DCSE work.

6.7.9 A New CASTEP and ONETEP Geometry Optimiser

This is the fourth DCSE project for the development of CASTEP, and was supervised by Dr Matt Probert of the Department of Physics at the University of York, who is the chair of the UK Car-Parrinello Consortium (UKCP). This work is also applicable to the ONETEP (Order-N Electronic Total Energy Package). The project was allocated 12 months support at the September 2009 round. NAG's Jolyon Aarons performed the work between September 2010 and September 2011.

The aim of the project was to reduce the memory-per-core requirements of the codes and thus improve the memory scaling of the geometry optimisation parts. Prior to this work both codes (since they have a common heritage and therefore similar implementations) used the Broyden-Fletcher-Goldfarb-Shanno (BFGS) method for non-linear optimisation. The BFGS method scales as $O(N^2)$, which limited the number of atoms of a simulation. For the inverse Hessian matrix SVD, this was replaced with the Limited-BFGS (L-BFGS) algorithm which provides $O(N)$ scaling for both codes.

New calculations on systems involving tens of thousands of atoms are now possible and general geometry optimisations can utilise more cores per HECToR node. The developments have gone into the main codebase of CASTEP and been passed back to the ONETEP developers.

6.7.10 Multi-core performance and domain choice for DL_POLY_3

The third DL_POLY DCSE project was supervised by Dr Ilian Todorov of STFC. This project was allocated 10 months support at the September 2009 round. The aims of the project were to optimise the multi-core performance by updating routines for the link cell, ewald and constraint force calculations, and to enable the use of processor counts other than powers of 2 by removing a restriction on the FFT part of the code.

Valène Pellissier of NAG performed the multi-core work between January 2010 and September 2010 and Ian Bush, also from NAG, worked on the FFT grid enhancements between May 2010 and July 2011. All of these developments have been introduced within the DL_POLY_4 code base (which is now the successor to DL_POLY_3). The multi-core optimisation work has resulted in roughly a 10-25% increase in performance on Phase 2b, depending upon the number of cores. DL_POLY_4 is no longer restricted to using 2^n cores, which also allows the scientist to study a system of interest without having to artificially inflate it to fit the restrictions of the code.

6.7.11 Improving Load Balancing and Parallel Partitioning in ICOM

Fluidity (ICOM) is an open-source, multi-scale, general-purpose CFD model, developed by the Applied Modelling and Computation Group (AMCG) at Imperial College London. This project was supervised by Dr Jon Hill of AMCG. It was allocated 12 months effort at the September 2009 round. Paul Woodhams of NAG performed the work between November 2009 and October 2011.

The main objective of this work was to improve the scalability of Fluidity by improved load balancing through the use of the Zoltan library, a collection of routines for parallel partitioning, load balancing and data management. This package was implemented successfully and has now been made the default solution for the Fluidity trunk, development branches and releases. In addition to the scaling performance of Fluidity, the maintainability, extensibility and functionality of the repartitioning solution has also been improved. The new Zoltan-based implementation has already been extended to allow parallel periodic problems to be solved by Fluidity. Detectors have also been implemented allowing them to be used within all parallel adaptive simulations. Another benefit of the Zoltan implementation is that it allows any element type to be used in parallel adaptive simulations. This will allow Fluidity to be used for new science.

The centrally installed version of Fluidity consumed 9.5M AUs on Phase 2a Between October 2009 and June 2011, and 0.6M AUs on Phase 2b between December 2010 and October 2011.

6.7.12 Optimising the performance of the VASP code

The VASP code is currently the single most used application on HECToR, used for performing *ab initio* quantum-mechanical molecular dynamics. This DCSE project was proposed by Professor Richard Catlow of the Department of Chemistry at UCL, and was allocated 12 months support at the September 2009 round.

The aims of the project were to optimise the parallel performance of VASP 5.2.2 for a set of Materials Chemistry Consortium (MCC) recommended benchmarks and also to implement k-point parallelism. Andrew Turner of EPCC performed the optimisation from January 2010 to August 2010 and Asimina Maniopoulou of NAG performed the k-point work from February 2010 to July 2011. For the non k-point work, recommended settings for NPAR were documented for five MCC benchmark cases. The code was then optimised by the implementation of global collectives and a shared-memory approach to aggregate intra node messages. This work improved scalability to more than 128 cores with up to 12 times speed ups on Phase 2b for some systems. An extra level of parallelism was also implemented by the addition of k-point parallelisation, whereby a group of cores may now be configured to perform calculations for a subset of k-points. Although k-point parallelism is not applicable to all systems, there are many cases where parallelisation over k-points is a useful technique; in such cases a speed-up of at least 2 times relative to the original code on the same core count can be achieved. In particular, for a unit cell of Litharge (α-PbO) with a total of 4 atoms using 108 k-points, the k-point code runs more than 5 times quicker than the original when using more than 1600 cores.

The centrally installed version of VASP used 187.8M AUs on Phase 2a between October 2009 and June 2011, and 191.8M AUs on Phase 2b between December 2010 and October 2011. The developments from this project are available on HECToR for all VASP users.

6.7.13 Implementation of a Divide and Conquer Strategy for the Materials Modelling code CRYSTAL

The second DCSE for CRYSTAL was supervised by Dr Barbara Montanari of STFC. This project was allocated 18 months effort at the September 2009 round. Daniel Jones of NAG performed the work between August 2010 and November 2011.

A divide and conquer algorithm to compute the energy of a system using density functional theory was implemented in CRYSTAL. The time taken to compute the density matrix using the divide and conquer algorithm now scales almost linearly with system size. While more work would be needed to be able to apply this algorithm to extended covalent networks and ionic systems, the method developed can be used for the *ab initio* study of weakly interacting systems. Project task farming was added to the code and performance showing convergence of the self-consistent fields for all subsystems also shows near perfect linear scaling. A method to improve the initial guess of the density matrix was also developed, which may be particularly useful for Molecular crystal, Macromolecule and Supercell simulations.

These developments are currently being used by the community of CRYSTAL users on HECToR, in particular at Imperial College London for the investigation of phase diagrams for varying chemical potentials.

6.7.14 Developing the TELEMAC system for HECToR

The TELEMAC-2D unstructured coastal modelling code is a new application on HECToR. TELEMAC uses the Finite Element Method to simulate free-surface flows, mainly environmental, dealing with current distribution, sediment transport, wave propagation, flooding and/or water quality treatment. The aim of this DCSE project was to improve the performance of TELEMAC on many-core architectures by the use of OpenMP within TELEMAC's underlying finite element library, BIEF. It was hoped that this would enable increased usage of TELEMAC on HECToR.

The project was allocated 12 months support at the January 2010 round and was supervised by Professor David Emerson of STFC. Zhi Shang, also of STFC, performed the work between October 2010 and September 2011. OpenMP directives were implemented in the Conjugate Gradient (CG) solver and in the Positive Streamwise Invariant (PSI) scheme; in total 200 subroutines were updated. Benchmarking tests demonstrated that a simulation using MPI / OpenMP (384 cores and 2 threads) in the Conjugate Gradient solver is now 80% faster than a simulation using pure MPI with the same number of cores on Phase 2b.

6.7.15 SPRINTing further with HECToR

SPRINT is an add-on package for the R statistical application. This DCSE project was allocated 6 months effort at the January 2010 round. The work was supervised by Terry Sloan of EPCC and performed by Lawrence Mitchell, also of EPCC, between October 2010 and April 2011. The aim of the project was to build on the success of the previous SPRINT DCSE, which had resulted in parallel implementations for the Pearson correlation function and the permutation testing function. This further work was to implement two more statistical functions, namely, the randomForest decision tree classifier and rank product.

A parallel wrapper was added around the serial randomForest algorithm along with a tree reduction approach for combining results in parallel. Typically, a 40 times speed up can now be achieved. A task parallel method using the existing serial rank product calculation was also implemented by distributing the bootstrap samples. For certain problem sizes, excellent scalability has been shown on Phase 2b for more than 512 cores. This work has also been presented at several international workshops and conferences and the new code is in the current release of SPRINT. Researchers at the Division of Pathway Medicine at the University of Edinburgh are now able to perform analyses in minutes instead of days. So far, around 629K AUs of the b10 project allocation has been consumed by SPRINT.

6.7.16 Microstructurally Faithful Modelling Of Materials

ParaFEM is a portable library of subroutines for parallel finite element analysis. This DCSE project was allocated 9 months support at the June 2010 round. Dr Lee Margetts of the University of Manchester supervised the work which was performed by Louise Lever, also from the University of Manchester, from December 2010 to November 2011.

The aim of this project was to widely extend the range of problems that could be solved with ParaFEM on HECToR, and to enable improved microstructurally-faithful models of materials. Several new user-defined material property subroutines (UMATs) have been implemented and optimised. The UMATs are Fortran 90 subroutines that implement a specific material behaviour such as linear elasticity or strain hardening plasticity. To further enable the visualisation capabilities of ParaFEM, the ParaFEM-Viewer was also ported to use the ParaView application, for performing parallel I/O on large datasets. These developments will be introduced within the main ParaFEM code base and are available to users on HECToR.

6.7.17 HECToR enabled Step Change in Turbulent Multiphase Combustion Simulations

The Direct Simulation of Turbulence and Combustion (DSTAR) code is used to conduct direct numerical simulation (DNS) of turbulent multiphase combustion. This DCSE project was allocated 6 months effort at the June 2010 round. Professor Kai Luo of the School of Engineering and the Environment at the University of Southampton supervised the work which was performed by Lucian Anton of NAG, between September 2010 and March 2011. The aim of the project was to improve the scalability of DSTAR in order to enable the code to use HECToR to conduct the world’s largest DNS of turbulent multiphase combustion. An improved domain decomposition was implemented by using the 2DECOMP&FFT library (which had already been developed for the Incompact3D DCSE project). The new code achieved approximately 50% scaling efficiency on Phase 2b for cubic grids of size 768 to 1536, with up to 18,432 MPI tasks. The original ASCII I/O method of using many processes to each write their own local monitoring/visualisation data was replaced with a significantly faster approach which uses a dedicated process perform the I/O. Also, the main Fortran 77 source code was re-developed as MODULES, and loops were reordered to improve cache memory usage, giving approximately a 20% increase in performance.

The improved DSTAR can now utilise approximately 50 times more MPI tasks and grid points. DSTAR has been used for the project “HECToR-Enabled Step Change in Turbulent Multiphase Combustion Simulation” (EP/I000801/1, 09/2010 – 12/2011), e186 with 32M

AUs. There are plans to use DSTAR in further projects on HECToR for Simulations of Thermal and Reactive Flow and Simulations of Combustion Instability.

6.7.18 Optimization of the MPI parallel RMT code for HECToR and likely successors

The R-Matrix with time-dependence (RMT) code was developed as part of the HELIUM DCSE project. It is the primary research tool of a collaboration between groups at The Open University, UCL, and project PI, Professor Ken Taylor of the Department of Applied Mathematics & Theoretical Physics at QUB. This DCSE project was allocated 6 months of effort at the December 2010 round. Jonathan Parker, also of QUB, performed the work during April to September 2011.

The aims of the project were to optimise the load balancing and communications for the RMT Inner Region (IR) with the RMT Outer Region (OR) for up to 10,000 cores on HECToR, and enhance the time-propagator in order to improve the efficiency of the numerical integration. Optimal parameters for balanced computation between the IR and OR were determined. A Red-Black load-balancing method was implemented for the IR/OR communications, this now divides the IR into two sets of cores which communicate independently with the OR. Computation was also moved from the OR to the IR (thereby reducing the information exchanged between the regions each time-step), and a single dedicated core has been assigned to the inter-region communication giving a 15-30% speed up. For the new time-propagator, a finite-difference algorithm based on a least-squares polynomial fit to the wavefunction was developed and implemented. A time step 1.8 times larger than was previously possible can now be used. This has enabled an 80% improvement in integration speed, while still preserving the stability and accuracy of the integration. Together with the use of a higher order propagator, a test case using 8,192 cores (IR+OR) demonstrated an increase in the integration speed by a factor of 1.7.

6.8 Ongoing DCSE Projects

6.8.1 Enhancing Conquest for accurate, scalable simulation of entire biological molecules

This is the second Conquest DCSE project, and is supervised by Dr David Bowler of the Department of Physics and Astronomy at UCL. The project was allocated 12 months support at the June 2010 round. Lianheng Tong, also of UCL, began the work February 2011 and is scheduled to complete in January 2012. The objectives are to:

- Develop an efficient O(N), parallel, linear scaling implementation of the van der Waals density functional in Conquest.
- Implement spin polarisation within Conquest.

6.8.2 Combined-Multicore Parallelism for the UK electron-atom scattering Inner Region R-matrix codes on HECToR

Following on from a previous DCSE project which targeted the Outer Region R-matrix code, PRMAT, this work will enhance the Inner Region codes for use within the UK-RAMP consortium. The project was allocated 12 months effort at the June 2010 round and is being supervised by Dr Martin Plummer of STFC. Andrew Sunderland and Cliff Noble, also of STFC, started the work in April 2011 and are scheduled to complete in March 2012. The common objectives are to upgrade the RAD, ANG and HAM(PDG) codes for improved OpenMP and task management, specific objectives are:

- Enhanced continuum function generation in RAD.
- Addition of the 'xstream' parallel I/O library to ANG.
- Improved dipole matrix calculations in HAM.

6.8.3 Optimisation of VASP Density Functional Theory (DFT)/Hartree-Fock (HF) hybrid functional code using mixed-mode parallelism

This project was allocated 12 months support at the June 2010 round and is the second DCSE project to improve the performance of VASP on HECToR. Dr Scott Woodley of the Department of Chemistry at UCL is supervising the work, which is aimed at enhancing the hybrid-DFT performance of the code. This will enable VASP to efficiently utilise more than 256 cores on HECToR. Andy Turner of EPCC started working on the project in November 2010, and Gavin Pringle, also of EPCC, took over the work in October 2011. The technical objectives are to add OpenMP parallelisation to the Projector-Augmented Wave routines and the hybrid DFT routines. The work is scheduled to be completed in October 2012.

Boosting the scaling performance of CASTEP: enabling next generation HPC for next generation science

This is the fourth CASTEP DCSE project, and was proposed by Dr Keith Refson of STFC. The project was allocated 10 months effort at the June 2010 round. The objectives are to improve:

- the MPI Buffer memory and optimise the I/O;
- the band-parallel capability of CASTEP for multi-core.

Dominik Jochym, also of STFC, performed the MPI buffer and I/O work between April 2011 and November 2011. Jolyon Aarons from the Department of Physics at the University of York started work on improving the band-parallel capability of CASTEP in October 2011, and is scheduled to complete the work in January 2012.

6.8.4 CP2K - Sparse Linear Algebra on 1000s of cores

Following on from two previous DCSE projects for the CP2K atomistic and molecular simulation code, this work will improve the sparse matrix (DBCSR) library performance in CP2K. The previous projects have targeted load balancing and OpenMP/MPI parallelism. This project was allocated 6 months support at the June 2010 round, and is being supervised by Dr Ben Slater of the Department of Chemistry at UCL. Iain Bethune of EPCC started the work in January 2011 and is scheduled to complete in January 2012. The objectives are:

- Implementation of a reordered 2D process distribution, automatic message type selection, and row/column communicators in DBCSR.
- Improved load balancing and on-node performance for the key multiplication routines.
- Implementation of automatic selection of optimal communication method and higher OpenMP efficiency.

The previous improvements to CP2K made with DCSE support have been introduced into the main code base and are in use on HECToR. Between October 2009 and June 2011 the centrally installed version of CP2K used 49.6M AUs on Phase 2a, and between December 2010 and October 2011 it used 121M AUs on Phase 2b.

6.8.5 Developing hybrid OpenMP/MPI parallelism for Fluidity/ICOM next generation geophysical fluid modelling technology

This is the third Fluidity (ICOM) DCSE project. The overall aim of this work is to enhance the multi-core performance of the code, and thus to enable increased scalability and efficiency on HECToR. The project was allocated 12 months effort at the June 2010 round. Dr Andrew Sunderland of STFC is supervising the work and Xiaohu Guo, also of STFC, began working on the project in October 2010 and is scheduled to complete in September 2012. The objectives are:

- Implementation of MPI/OpenMP mixed-mode parallelisation of the finite element assembly stage in Fluidity.
- Optimization of the HYPRE Library usage for linear preconditioners/solvers for large core counts.
- Benchmarking and code re-engineering for hybrid mesh adaptivity.

6.8.6 Developing NEMO for Large Multi-core Scalar Systems

This project was allocated 12 months support at the June 2010 round and is the second DCSE project to improve the performance of the NEMO Ocean modelling application on HECToR. Dr Stephen Pickles of STFC is both supervising and working on the project, along with Dr Andrew Porter, also of STFC. Dr Pickles worked on the project between January and December 2011. Dr Porter will take over the work in January 2012 and is scheduled to complete the project in June 2012. The objectives are:

- Provision of an ensemble capability in NEMO.
- Multicore-aware partitioning and halo exchange optimisation.
- Array index re-ordering.

6.8.7 Parallelisation and porting of UKRMol-in the electron molecule scattering inner region R-matrix codes

The UKRMol-in code is another application which is used within the UK-RAMP consortium. This DCSE project was proposed by Dr Jimena Gorfinkiel of the Department of Physics & Astronomy at the Open University. It was allocated 12 months effort at the December 2010 round. The objectives are:

- Parallelisation of the Hamiltonian construction.
- Optimization of a sparse-matrix vector multiplication OpenMP/MPI approach.
- Full implementation of an optimized diagonalisation step.

Michael Lysaght of the Open University, who had previously worked on the HELIUM DCSE project, carried out the parallel Hamiltonian construction implementation between April and November 2011. Paul Roberts of NAG took over in December 2011 and is scheduled to complete the remaining work on further optimization of the diagonalisation step in May 2012.

6.8.8 CABARET on Jet Flap Noise and Quasigeostrophic Ocean Circulation Models

This project was allocated 6 months support at the December 2010 round and is the second DCSE project concerning the CABARET application. The work was proposed by Dr Sergey Karabasov of the Department of Engineering at the University of Cambridge. The first objective of the project is to improve the performance of the existing unstructured version of the CABARET code, which was ported to HECToR in another DCSE project. The second objective is to port a scalable implementation of a 3D Quasigeostrophic(QG) ocean model, which uses CABARET, to HECToR. Phil Ridley of NAG began working on the project in August 2011 and is scheduled to finish in July 2012.

6.8.9 Scalable coupling of Molecular Dynamics and Direct Numerical Simulation of multi-scale flows

The overall objective of this project is to implement, on HECToR, a fully coupled and scalable simulation of molecular dynamics (MD) with CFD. Dr Tamer Zaki of the Department of Mechanical Engineering at Imperial College London proposed the work, which was approved for 6 months effort at the December 2010 round. As part of this work, a coupled continuum MD model for the TransFlow CFD solver and the Stream (MD) solver has been developed. The coupler includes a communications infrastructure (for internal data management) and interface module (for data access between applications). Lucian Anton of NAG began working on the project in April 2011 and is scheduled to finish in March 2012.

6.8.10 Direct Numerical Simulations (DNS) of turbulent fluid flows

This is the second DCSE project concerning the Incompact3D application. The objective of this work is to port two Incompact3D related applications to HECToR, namely Compact3d and QuasiCompact3D, and to implement a new communication library for particle tracking for efficient use on HECToR. This project was allocated 6 months support at the December 2010 round. Professor John C. Vassilicos of the Department of Aeronautics at Imperial College London is supervising the work, and Ning Li of NAG began working on the project in March 2011 and is scheduled to complete in February 2012.

6.8.11 Improved Data Distribution Routines for Gyrokinetic Plasma Simulations

The aim of this project is improve the method of indirect addressing for data access within the Gs2 code, which is used for modelling gyrokinetic flux turbulence in magnetised plasmas. A total of 150M AUs is allocated for use with Gs2 and related applications. Dr Colin Roach of the EURATOM/Culham Centre for Fusion Energy proposed the work, which was approved for 6 months of effort at the December 2010 round. Adrian Jackson of EPCC began working on the project in March 2011 and is scheduled to finish in January 2012.

6.8.12 RMT for High Harmonic Generation

This is the second DCSE project concerning the R-Matrix with time-dependence (RMT) code, which is also related to the UK-RAMP consortium project on HECToR. The aim of this project is to develop new code for the calculation of High Harmonic Generation (HHG) processes, adding further functionality to the scalable RMT code. The project was proposed by Professor Ken Taylor of the Department of Applied Mathematics & Theoretical Physics at QUB. The project was allocated 6 months of support at the June 2011 round. Jonathan Parker, also of QUB began working on the project in October 2011 and is scheduled to finish in March 2012.

6.8.13 Improved Scaling for Direct Numerical Simulations of Turbulence

The overall objective of this project is to improve the scaling performance of the turbulence applications SS3F and SWT on HECToR. This will be achieved by replacing the existing FFT subroutines with the FFTW library and developing routines for improved data decomposition. Professor Gary Coleman of the Department of Engineering and the Environment at the University of Southampton proposed the work, which was approved for 6 months of effort at the June 2011 round. Roderick Johnstone, also of the University of

Southampton, began work on the project in November 2011 and is scheduled to finish in July 2012.

6.8.14 Preparing DL_POLY_4 for the Exascale

This is the fourth DL_POLY DCSE project, and was proposed by Dr Ilian Todorov of STFC. The project was allocated 12 months support at the June 2011 round. The objectives are to implement a full mixed-mode OpenMP/MPI version to exploit the shared memory features of HECToR, and to enable billion atom simulations by implementing a 64-bit integer representation within the code. Ian Bush of NAG started the mixed-mode work in November 2011 and is scheduled to finish in October 2012. Asimina Maniopoulou, also of NAG, started work on the 64-bit integer representation in December 2011, and is scheduled to complete this part of the work in July 2012.

6.9 New DCSE Projects Starting 2012

6.9.1 Bootstrapping and support vector machines with R and SPRINT on HECToR

The overall aim of this third SPRINT DCSE project is to build upon the success of the previous work. The specific objectives of this project are to implement two R functions, bootstrapping and support vector machines (SVM), for use with SPRINT on HECToR. The addition of this functionality in a scalable implementation on HECToR is expected to bring wider use of SPRINT and HECToR for the bioscience community. The project was allocated 6 months effort at the June 2011 round and will be supervised by Mr Terry Sloan of EPCC. Michal Piotrowski, also of EPCC, is scheduled to begin the work in February with a scheduled completion in July 2012.

6.9.2 Improving CONQUEST for biomolecular simulations: scaling and molecular dynamics

This is the third Conquest DCSE project, and will be supervised by Dr David Bowler of the Department of Physics and Astronomy at UCL. The project was allocated 8 months support at the June 2011 round. Lianheng Tong, also of UCL, will begin the work in February with a scheduled completion in September 2012. The aim of the project is to develop a new MPI / OpenMP version of Conquest to enable efficient scalability on up to 4,096 cores on Phase 3.

6.9.3 Enhancement of a high-order CFD solver for many-core architecture

This project was proposed by Professor Paul Tucker of the Department of Engineering at the University of Cambridge and was allocated 4 months effort at the June 2011 round. The Block Overset Fast Flow Solver (BOFFS) is a structured LES model, capable of resolving the low frequency components in flow structures with lower computational overheads than conventional LES solvers. This new hybrid code has been developed by Professor Tucker's group and the aim of this project is to improve the efficiency of the inter-node communications, which will enable scalability to over 1000 cores on HECToR. The project will start in January 2012, with a scheduled completion in September 2012.

6.9.4 VOX-FE Large Scale FE Bone Modelling on HECToR

VOX-FE is a finite element solver specifically designed to solve linear elasticity problems relating to bone modelling. The simulation grid for VOX-FE is generated directly from microCT scan data which ensures an accurate representation, as each scan voxel becomes a finite element. VOX-FE has recently been ported to HECToR and the aim of this DCSE project is to enable it to utilise HECToR for an entire model of femur or bone/implant interface. The proposed work will involve improvement to the data format, I/O performance and load balancing. The project will be supervised by Professor Michael Fagan of the Department of Engineering at the University of Hull. Nick Jackson of EPCC will begin the work in February 2012 and is scheduled to complete in September 2012.

6.9.5 Micromagnetic modelling of naturally occurring magnetic mineral systems

This is the second DCSE project concerning the MicroMag application. The specific objective of this work is to develop the existing parallel implementation by improving the initialisation step and the load balancing for the unstructured data. This will then enable large-scale simulations with increased system size and precision and also facilitate more realistic theory to compare with experimental observation. The project was allocated 6 months effort at the June 2011 round and will be supervised by Professor Wyn Williams of the School of Geosciences at the University of Edinburgh. Chris Maynard of EPCC will begin the work in February 2012 and is scheduled to complete in November 2012.

Appendix A: Terminology

TSL : Threshold Service Level

FSL : Full Service Level

SDT : Scheduled Down Time

UDT : Unscheduled Down Time

WCT : Wall Clock Time

MTBF : Mean Time Between Failures = 732/Number of Failures

SP : Service Provision

SP Serviceability% = 100*(WCT-SDT-UDT(SP))/(WCT-SDT)

Technology Reliability % = 100*(1-(UDT(Technology)/(WCT-SDT))

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 file systems 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 on HECToR

Code	Project Title	Funding Body	Class	Principal Investigator	AUs allocated	AUs used	AUs left
EPSRC Projects (Class 1a listed first, followed by Class 1b, Class 2a, and Class 2b)							
c01	Support of EPSRC/STFC SLA	EPSRC	Class1a	Dr Richard Blake	50,803.70	37,126.50	13,677.10
e01	UK Turbulence Consortium	EPSRC	Class1a	Dr Gary N Coleman	483,969.90	80,691.10	403,141.50
e05	Materials Chemistry HPC Consortium	EPSRC	Class1a	Prof C Richard A Catlow	1,139,124	311,568.60	826,879.90
e10	GENIUS	EPSRC	Class1a	Prof Peter Coveney	257,748.20	9,829.20	247,919
e104	Fluid-Mechanical Models applied to Heart Failure	EPSRC	Class1a	Dr Nicolas Smiths	30,400	7,020.10	23,379.90
e105	Joint Euler/Lagrange Method for Multi-Scale Problems	EPSRC	Class1a	Dr Andreas M Kempf	1,300	297.3	1,002.70
e106	Numerical Simulation of Multiphase Flow: From Mesocales to	EPSRC	Class1a	Prof Kai Luo	3,650	0	3,650
e107	Parallel Brain Surgery Simulation	EPSRC	Class1a	Dr Stephane P. A. Bordas	6,000	713.2	5,286.80
e108	Jet Flap Noise	EPSRC	Class1a	Dr Sergey Karabasov	49,684.50	14,546.30	35,138.20
e110	Computational Aeroacoustics Consortium	EPSRC	Class1a	Prof Paul Tucker	140,110.30	58,022	82,026.30
e121	[dCSE] Improving Performance using Wannier functions	EPSRC	Class1a	Prof Maria Merlyne DeSouza	2,680.30	2,299.60	380.7
e122	Multiscale Modelling of Magnetised Plasma Turbulence	EPSRC	Class1a	Dr Colin M Roach	65,000	35,985.30	28,933.90
e124	Compressible Axisymmetric Flows	EPSRC	Class1a	Dr Richard D Sandberg	22,887.90	7,947.50	14,936.40
e125	Full configuration interaction quantum monte carlo	EPSRC	Class1a	Dr Ali Alavi	168,324.80	13,576.50	154,638.30

Code	Project Title	Funding Body	Class	Principal Investigator	AUs allocated	AUs used	AUs left
e126	Clean Coal Combustion: Burning Issues of Syngas Burning	EPSRC	Class1a	Prof Xi Jiang	25,584	8,271.40	17,312.60
e127	Alternative drag-reduction strategies	EPSRC	Class1a	Prof Michael Leschziner	7,000	1,167.90	5,832.10
e128	Rate-Controlled Constrained Equilibrium	EPSRC	Class1a	Dr Stelios Rigopoulos	7,092.10	3,494	3,598.10
e129	Novel Hybrid LES-RANS schemes [ICL]	EPSRC	Class1a	Prof Michael Leschziner	7,500	1,076.20	6,423.80
e130	Novel hybrid LES-RANS schemes [MAN]	EPSRC	Class1a	Prof Dominique Laurence	10,500	1,945.80	8,554.20
e141	A numerical study of turbulent manoeuvering-body wakes	EPSRC	Class1a	Dr Gary N Coleman	16,350	3,401.50	12,948.50
e143	Numerical Investigation of Jet Noise	EPSRC	Class1a	Dr Anurag Agarwal	0	0	0
e144	Numerical Simulation of Rotating Stall and Surge	EPSRC	Class1a	Dr Mehdi Vahdati	1,266	0.3	1,265.70
e145	UK-SHEC Consortium	EPSRC	Class1a	Dr T.J. Mays	1,191.90	367.8	821.6
e149	Fractal-generated turbulence and mixing: flow physics and	EPSRC	Class1a	Prof Christos Vassilicos	68,082.50	46,650.60	21,431.90
e155	Modelling Cholesterol Deposits	EPSRC	Class1a	Dr David Quigley	10,000	161.7	9,838.30
e158	Novel Asynchronous Algorithms	EPSRC	Class1a	Prof Nicholas J Higham	500	279.1	220.9
e159	Multi-layered Abstractions for PDEs	EPSRC	Class1a	Prof Paul Kelly	3,816	11.8	3,804.20
e160	Sustainable Software Generation Tools	EPSRC	Class1a	Prof Paul Kelly	20,208.10	0.9	20,207.10
e161	Properties and Dynamics of Atomic Bose-Einstein Condensates	EPSRC	Class1a	Dr A White	69,895.50	0	69,895.50
e165	Multi-scale simulation of intense laser plasma interactions	EPSRC	Class1a	Dr Tony Arber	4,872	0	4,872
e175	Fine-Scale Turbulence	EPSRC	Class1a	Dr Richard D Sandberg	50,000	509.4	49,334.90

Code	Project Title	Funding Body	Class	Principal Investigator	AUs allocated	AUs used	AUs left
e179	Non-conservative dynamics	EPSRC	Class1a	Dr Daniel Dundas	87,000	705.7	86,294.30
e182	Advanced Modelling of Two-Phase Reacting Flow	EPSRC	Class1a	Dr Edward S Richardson	8,150.20	0	8,150.20
e183	Analysis of Processes in Hydrocarbon Fuel Droplets	EPSRC	Class1a	Prof Sergei Sazhin	8,640	0	8,640
e184	UK-RAMP	EPSRC	Class1a	Prof Ken Taylor	130,500	732	129,768
e185	Chemistry of ceramic materials	EPSRC	Class1a	Prof John Harding	340,000	6,033.10	333,966.90
e186	Step Change in Combustion Simulation	EPSRC	Class1a	Prof Kai Luo	40,000	18,772.50	21,172.40
e187	IAGP: Integrated Assessment of Geoengineering Proposals	EPSRC	Class1a	Prof Piers Fosters	6,030.20	4.8	6,025.40
e191	CFD Analysis of Flight Dynamics	EPSRC	Class1a	Prof Kenneth Badcock	40,500	4,413.10	36,086.90
e202	Quantum Monte Carlo simulations	EPSRC	Class1a	Prof Matthew Foulkes	38,345	0	38,345
e203	BeatBox - Realistic Cardiac Simulations	EPSRC	Class1a	Prof Vadim Biktashev	4,400	50.7	4,349.30
e204	Rare Events via Parallel Forward Flux Sampling	EPSRC	Class1a	Dr Rosalind Allen	5,000	0	5,000
e206	FLAME Agent-Based Simulation Framework	EPSRC	Class1a	Prof Christopher Greenough	410	0	410
e207	Developing DL_POLY Molecular Dynamics Simulation code	EPSRC	Class1a	Dr Kostya Trachenko	25,857.60	0	25,857.60
e211	Dendrite simulation	EPSRC	Class1a	Dr Jiawei Mi	300	1.1	298.9
e226	Novel Vibrational Spectroscopic Techniques	EPSRC	Class1a	Dr Andrew D Burnett	1,032.30	0	1,032.30
e228	Development of the potential of doped metal-oxide nanotubes	EPSRC	Class1a	Dr Gilberto Teobaldi	4,918.30	153.8	4,764.50
e229	DTC in Complex Systems Simulations	EPSRC	Class1a	Prof Jonathan W Essex	50,000	0	50,000

Code	Project Title	Funding Body	Class	Principal Investigator	AUs allocated	AUs used	AUs left
e241	Potential Energy Surfaces for Bio-molecular Simulations	EPSRC	Class1a	Dr Lorna Smith	500	0	500
e42	Computational Combustion for Engineering Applications	EPSRC	Class1a	Prof Kai Luo	32,000	30,171.30	1,828.70
e63	UK Applied Aerodynamics Consortium 2	EPSRC	Class1a	Dr Nick Hills	30,925.30	31,172.70	-247.4
e68	Hydrogenation Reactions at Metal Surfaces	EPSRC	Class1a	Prof. Angelos Michaelides	50,000	49,791.10	208.9
e71	Simulating the control of calcite crystallisation	EPSRC	Class1a	Prof John Harding	130,403.50	49,479.60	80,912.30
e76	HELIUM Developments	EPSRC	Class1a	Prof Ken Taylor	42,521.80	34,613.20	7,908.50
e84	Vortical Mode Interactions	EPSRC	Class1a	Dr Tamer Zaki	9,600	3,203.10	6,396.90
e85	Study of Interacting Turbulent Flames	EPSRC	Class1a	Dr N Swaminathan	8,088.60	3,763.70	4,324.90
e89	Support for UK Car-Parrinello Consortium	EPSRC	Class1a	Dr Matt Probert	360,100	262,295.30	96,860.70
e92	Dynamo Action In Compressible Convection	EPSRC	Class1a	Mr Paul Bushby	4,075	4,074.40	0.6
j01	JST	EPSRC	Class1a	Dr Andrew R Turner	71,990.70	16,059.20	55,838.70
e139	Scalability Optimization for Largescale in-silico Simulations	EPSRC	Class1b	Dr Gernot Plank	3,121.10	588.9	2,532.20
e173	Performance of oomph-lib in largescale parallel computations	EPSRC	Class1b	Prof Matthias Heil	4,800	245.1	4,554.90
e174	3D instabilities in two-layer flows	EPSRC	Class1b	Dr Prashant Valluri	9,243.40	551.8	8,691.60
e177	Amorphous structures of mirror coatings	EPSRC	Class1b	Dr Ian Maclaren	5,700.80	301	5,399.70
e193	Colloids in Cholesteric Liquid Crystals	EPSRC	Class1b	Dr Davide Marenduzzo	28,793.90	15,035.50	13,642.90
e205	Feasilibility study of fine sediment transport	EPSRC	Class1b	Dr Ming Li	3,000	129.6	2,870.40

Code	Project Title	Funding Body	Class	Principal Investigator	AUs allocated	AUs used	AUs left
e214	MD Studies of Low Salinity Enhanced Oil Recovery Mechanisms	EPSRC	Class1b	Prof Peter Coveney	3,086.60	0	3,086.60
e215	GIPAW DFT Calculation of NMR Parameters in Rare Earth Materials	EPSRC	Class1b	Dr John V Hanna	8,170	2,962.60	5,178.40
e216	Self-organised Lipid layers on Mercury	EPSRC	Class1b	Dr Pietro Ballone	1,535	693.8	800
e217	Exploring a Conformational Switch in a Macromolecule	EPSRC	Class1b	Dr Philip Biggin	2,835.40	879.9	1,955.60
e218	Computational Electron Collison Experiments using 2DRMP	EPSRC	Class1b	Dr Penny Scott	1,449.60	20.7	1,428.90
e219	Gust generation modelling for aeronautical purposes	EPSRC	Class1b	Prof Oubay Hassan	1,620	1,108.50	511.5
e220	Study of interacting turbulent flames 2	EPSRC	Class1b	Dr N Swaminathan	16,920	0	16,920
e233	Lengthscale bridging of biophysical systems	EPSRC	Class1b	Prof Jason Crain	10,400.60	229.5	10,171.10
e234	Simulations of carbon electrodes with ionic electrolytes	EPSRC	Class1b	Prof. Paul A Madden	1,968.50	0	1,968.50
e82	ONETEP: linear-scaling method on High Performance Computers	EPSRC	Class1b	Dr Peter Haynes	1,105.40	866.8	238.5
e210	The Defect Chemistry of TiO2	EPSRC	Class2a	Prof Russell Howe	300	177.3	122.7
e213	Condensation/Evaporation Heat Transfer in Micro/Nanochannels	EPSRC	Class2a	Dr Huasheng Wang	400	0	400
e222	Integrated Drug Delivery Systems	EPSRC	Class2a	Dr Charles Laughton	400	430.9	-30.9
e223	Numerical modelling of aorta dissection	EPSRC	Class2a	Prof. Xiaoyu Luo	300	0	300
e224	Electronic properties of inorganic-organic hybrid materials	EPSRC	Class2a	Prof Anthony K Cheetham	400	36.1	363.9
e225	New Ru and Ir Chromophores for Solar Cell Devices	EPSRC	Class2a	Dr Paul Elliott	300	89.2	210.8
e227	OPL	EPSRC	Class2a	Dr Radhika R. S. Saksena	50	46.4	3.6

Code	Project Title	Funding Body	Class	Principal Investigator	AUs allocated	AUs used	AUs left
e230	Adsorption and Diffusion in Metal-Organic Frameworks	EPSRC	Class2a	Dr Ahmet Ozgur Yazaydin	400	163.8	223.5
e231	Rapid Alloy Solidification	EPSRC	Class2a	Prof Peter Jimack	400	0	400
e232	Flow field analysis around flap type wave energy devices	EPSRC	Class2a	Dr Matthew Folley	289.9	0	289.9
e235	Modelling offshore wind	EPSRC	Class2a	Prof Simon Watson	400	0	400
e236	Simulations of Optical Communications Systems	EPSRC	Class2a	Dr Marc Eberhard	400	0	400
e237	Simulating Coupled Protein Folding and Nucleic Acid Binding	EPSRC	Class2a	Dr Christopher Baker	400	0	400
e238	Porting to CAF and Experiments on the Peppermint Application	EPSRC	Class2a	Dr Stephen Jarvis	400	0	400
e239	Optimum Collection and Conversion of Light into Energy	EPSRC	Class2a	Dr Robert Paton	400	0	400
e242	Study of the Green Fluorescent Protein Fluorophore	EPSRC	Class2a	Dr Garth Jones	400	0	400
e156	Metal Conquest: efficient simulation of metals on petaflop	EPSRC	Class2b	Dr David Bowler	1,600	56.7	1,543.30
e240	MicroMag	EPSRC	Class2b	Prof Wyn Williams	800	2.4	797.6
STFC Projects							
p01	Atomic Physics for APARC	STFC	Class1a	Dr Penny Scott	10,002.70	666.3	9,336.40
NERC Projects							
n01	Global Ocean Modelling Consortium	NERC	Class1a	Dr Andrew C Coward	156,545.50	114,578.30	29,650.80
n02	NCAS (National Centre for Atmospheric Science)	NERC	Class1a	dr grenville gms lister	500,832.30	372,726.20	127,912.40
n03	Computational Mineral Physics Consortium	NERC	Class1a	Prof John P Brodholt	405,647	323,838.60	81,140.40

Code	Project Title	Funding Body	Class	Principal Investigator	AUs allocated	AUs used	AUs left
n04	Shelf Seas Consortium	NERC	Class1a	Dr Roger Proctor	104,161.50	80,595	23,550.80
u01	Melting of MgSiO3 Perovskite	NERC	Early use	Prof John P Brodholt	11,000	11,018.40	-18.4
BBSRC Projects							
b08	Int BioSim	BBSRC	Class1a	Mr Mark M Sansom	866	910	-44
b09	Circadian Clock	BBSRC	Class1a	Prof Andrew A Millar	2,000	1,393.90	606.1
b100	Widening the BBSRC HPC User Base	BBSRC	Class1a	Dr Michael Ball	10,000	632.5	9,367.50
b12	Flu Analysis on HECToR	BBSRC	Class1a	Mr Adrian Jackson	50	0	50
b13	Linear Scaling DFT for Biochemistry Applications	BBSRC	Class1a	Dr David Bowler	5,587.20	105.6	5,481.60
b14	Understanding supercoiling-dependent DNA recognition	BBSRC	Class1a	Prof Anthony Maxwell	42,600	0	42,600
Director's Time							
d11	NAIS	Directors Time	Service	Prof Mark Ainsworth	10,000	1,221.70	8,778.30
d15	HPC-GAP	Directors Time	Service	Dr David Henty	102	2.7	99.3
d16	ETC	Directors Time	Service	Dr Lorna Smith	501	199.6	301.4
d19	OpenFOAM Demo	Directors Time	Service	Dr Alan Gray	1,950	1,894.80	55.2
d21	GADGET	Directors Time	Service	Dr Adrian Jenkins	1,000	18.6	981.4
d23	TEXT FP7	Directors Time	Service	Dr Mark Bull	1,500	30.5	1,469.50
d24	SBSI	Directors Time	Service	Dr Stephen Gilmore	2,000	958.1	1,041.90

Code	Project Title	Funding Body	Class	Principal Investigator	AUs allocated	AUs used	AUs left
d25	Code Scaling	Directors Time	Service	Dr Ken Rice	51,500	6,571.20	44,928.80
d26	Guest Training Accounts	Directors Time	Service	Miss Elizabeth Sim	50	43.2	6.8
d27	RollsRoyce	Directors Time	Service	Mr Paul Graham	50	27.3	22.7
d29	Nu-FuSe	Directors Time	Service	Mr Adrian Jackson	500	0	500
d30	PARTRAC	Directors Time	Service	Dr Mark Sawyer	200	86.4	113.6
d31	Semileptonic Decay	Directors Time	Service	Prof Richard Kenway	1,000	0	1,000
d32	APOS-EU	Directors Time	Service	Dr Michele Weiland	1,000	124.6	875.4
d33	Mark Westwood's Project	Directors Time	Service	Mr Mark Westwood	100	8.9	91.1
d34	Msc 2011-2012	Directors Time	Service	Dr David Henty	1,000	17.7	982.3
d35	PhD	Directors Time	Service	Dr Mark Bull	10	0	10
d36	Genome	Directors Time	Service	Dr Alan Gray	3,460	0	3,460
d37	CRESTA	Directors Time	Service	Dr Lorna Smith	1,000	61.8	924.7
d38	Windfarm Simulation	Directors Time	Service	Mr Adrian Jackson	171	0	171
d39	NCSA access	Directors Time	Service	Mr Mark A Straka	1,000	52.1	947.9
x07	RSI	Directors Time	Service	Miss Elizabeth Sim	10	0	10
y09	Director's Time	Directors Time	Service	Prof Arthur S Trew	29,685.10	82.5	764.2
External Projects							

Code	Project Title	Funding Body	Class	Principal Investigator	AUs allocated	AUs used	AUs left
t01	NIMES: New Improved Muds from Environmental Sources.	External	Service	Dr Chris Greenwell	4,113.70	4,245.40	-131.8
x05	FIOS	External	Service	Mr Davy Virdee	1,130.10	1,076.60	53.5
e168	TEXT	External	Service	Dr Mark Bull	1,500	0	1,500
x01	HPC-Europa	External	Service	Dr Judy Hardy	25,564.80	15,928.40	9,608.80
x02	BlueArc (TDS)	External	Service	Mr M W Brown	1	0	1
x06	Rhymney	External	Service	Dr Mark Sawyer	4.5	0.1	4.4
PRACE Projects							
pr1u0701	EC4aPDEs-2	PRACE	Class1a	Dr Chris A Johnson	0	0	0
pr1u0702	HYDROGEN-ILs	PRACE	Class1a	Dr Chris A Johnson	0	0	0
pr1u0703	HELIXKINETICS	PRACE	Class1a	Dr Chris A Johnson	0	0	0
pr1u0704	HIFLY	PRACE	Class1a	Dr Chris A Johnson	0	0	0
pr1u0705	Tangrin	PRACE	Class1a	Dr Chris A Johnson	2,800	27.8	2,772.20
pr1u0706	SIVE-2	PRACE	Class1a	Dr Chris A Johnson	5,000	1,199.90	3,800.10
pr1u9999	PRACE Training	PRACE	Class1a	Dr David Henty	0	0	0