HECToR Annual Report 2012

01 January – 31 December 2012

Issue: Draft

Contents

1 Introduction

This report covers the period from 1 Jan 2012 at 08:00 to 1 Jan 2013 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/2012.php

2 Executive Summary

2012 was another positive year for the HECToR Service. The highlights of the service over the year included:

- Average utilisation on the XE6 in 2012 was 85% of *optimum*. Further details are available in Section 3.2 of the report.

- The upgrade to Phase 3 came to a conclusion in 1Q12. The remaining 10 cabinets of Interlagos were installed on 18th January, with acceptance tests completed in early February. The system is now a 30 cabinet Interlagos based system. This equates to 90,112 cores, with a peak capacity of over 800 T/Flops.

- The official Phase3 launch event took place in February. The event was attended by a number of senior stakeholders from the UK Research Councils, the University of Edinburgh and the government. The event was also attended by the winner of EPSRC's HECToR Cabinet Design competition.
- There were 12 technology-attributed service failures in 2012, as opposed to 10 in 2011. 5 of these were attributed to maintenance overruns as opposed to hardware failures. The overall MTBF was 732 hours compared to 586 hours in 2011. The XE6 has been very reliable since its introduction in 2011.
- A total of 4648 queries were handled in 2012 and the associated Helpdesk statistics for the year were excellent.

- A history of all user jobs and project details are maintained in the SAFE administration database. This proved invaluable in 2012 in supporting analyses into both the impact of utilisation on job backlog, and the spend profile of Class1a projects.

- HECToR was connected to the Research Data facility (RDF) in May. The RDF comprises of 7.8PB of disk storage, with 19.5PB of backup tape. Users from 10 EPSRC and 4 NERC consortia have access to the disk.

- At the request of a number of user groups, new queues supporting 24-hour jobs were introduced on 4th July for an initial 6 month trial period. The initial trial was deemed a success and the queues will now remain in place. Previously the longest job that could be supported outside a special reservation was 12 hours. These longer jobs have accounted for around 16% of the overall utilisation.

- The CSE training programme remains popular. To date the CSE have filled over 1600 places on training courses, and continue to receive positive feedback from attendees and invitations to return to institutions to provide further training in future.

- EPSRC has provided additional DCSE funding to continue the programme until September 2013 (the original service end date). In addition further funding has been provided to offer DCSE to non-HECToR users in a pilot programme that started in the summer. To date the DCSE programme has allocated over 65 person years of effort to improving codes deemed important to UK research. Further details on the 2012 completed and ongoing DCSE projects are available in Section 6 of the report.

- In support of the Phase3 upgrade the Cray Centre of Excellence (CoE) and Exascale Research Initiative (ERI) engaged in user training, workshops and high level support of users. The CoE was also engaged in code porting and optimisation activities. Examples of this are available in Section 5.2 of the report.

- A HECToR User Group (HUG) meeting was held in December in London.

3 Quantitative Metrics

3.1 Reliability

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

	1Q12	*2Q12*	*3Q12*	*4Q12*	*2012*
Incidents	60	28	45	54	187
Failures	5	3	2	2	12

The incidents above are primarily related to single node failures. Details on both the service failures and single node fails in 2012 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*	*1Q12*	*2Q12*	*3Q12*	*4Q12*	*2012*
Technology	Failures	5	3	2	2	12
	MTBF	439	732	1098	1098	732
Service Provision	Failures	0	0	0	0	0
	MTBF	∞	∞	∞	∞	∞
External/Other	Failures	0	0	0	0	0
	MTBF	∞	∞	∞	∞	∞
Overall	**Failures**	**5**	**3**	**2**	**2**	**12**
	MTBF	**439**	**732**	**1098**	**1098**	**732**

3.2 HECToR Utilisation

3.2.1 Phase 3 Utilisation

Overall utilisation of the XE6 in 2012 was 68%.

A history of all user jobs and project details are maintained in the SAFE administration database. This proved invaluable in 2012 in supporting analyses into both the impact of utilisation on job backlog, and the spend profile of Class1a projects. The chart below demonstrates the effect of utilisation on job backlog. The chart is based on SAFE accounting records dating back to August 2009. The utilisation on the service was compared to the waiting job backlog at that level of utilisation. As below the backlog has been seen to rise dramatically at utilisation levels above around 80%. This has now been defined as 'optimum' utilisation.

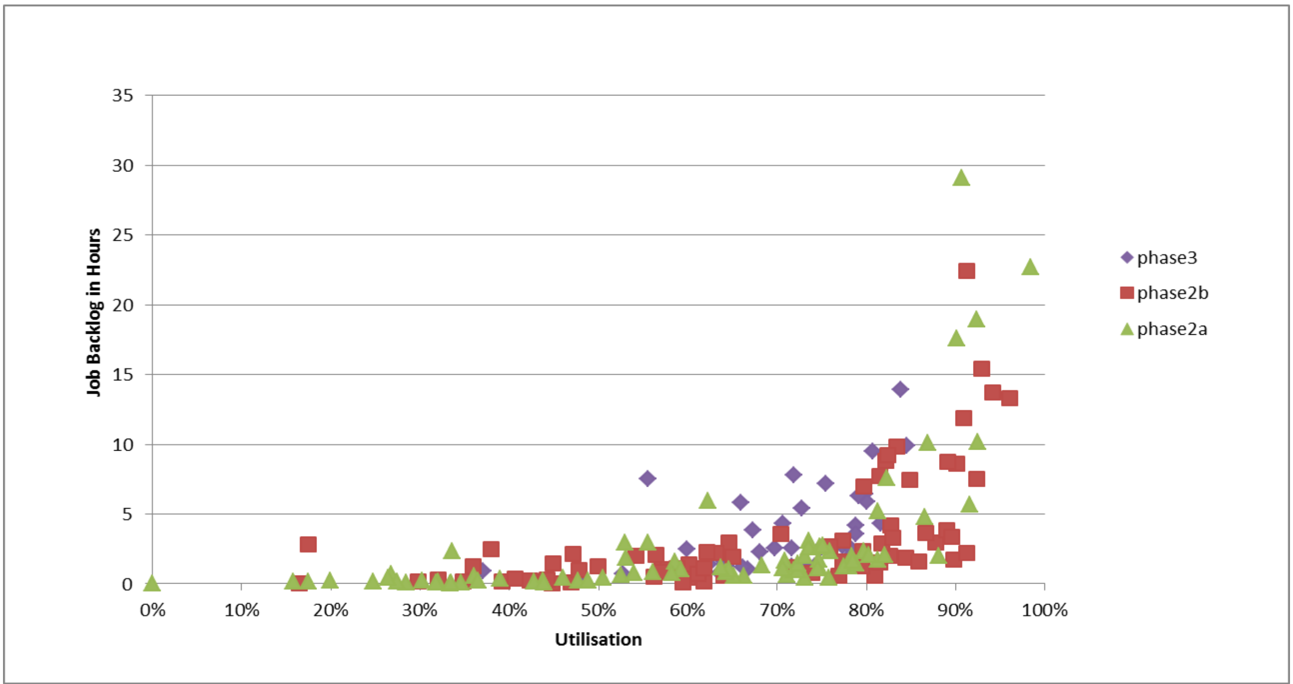

The graph below shows the monthly utilisation in 2012 relative to the 80% optimum value. Overall this equates to 85% across the year.

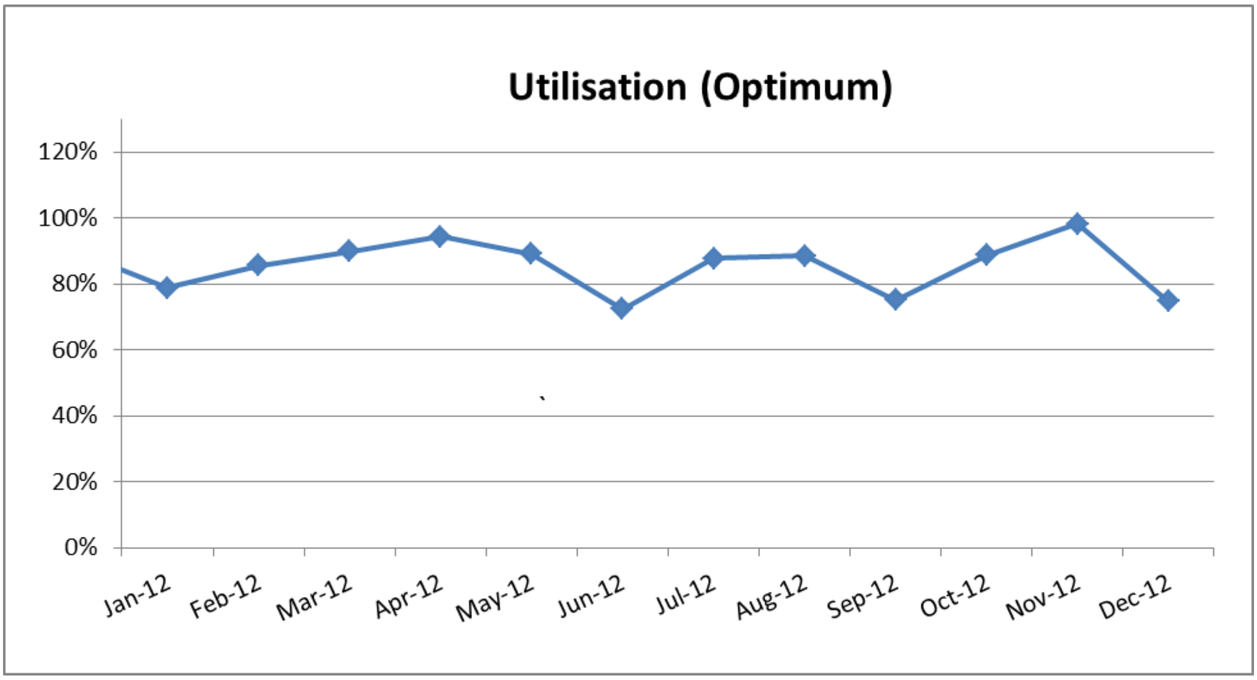

3.2.2 HECToR Utilisation by Project Class

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/Industrial Projects

The primary usage on HECToR is as expected from full peer reviewed 'Class 1a' projects; this accounted for 89% of the utilisation. Usage from Direct Access projects decreased from 7.7% in 2011 to 5% in 2012.

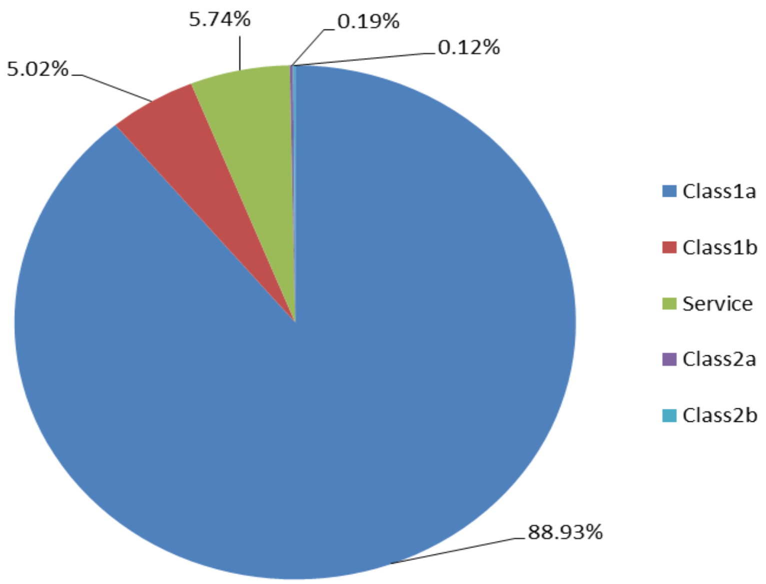

3.2.3 HECToR Utilisation by Consortium

As below, the main utilisation in 2012 came from the major EPSRC and NERC consortia as one would expect.

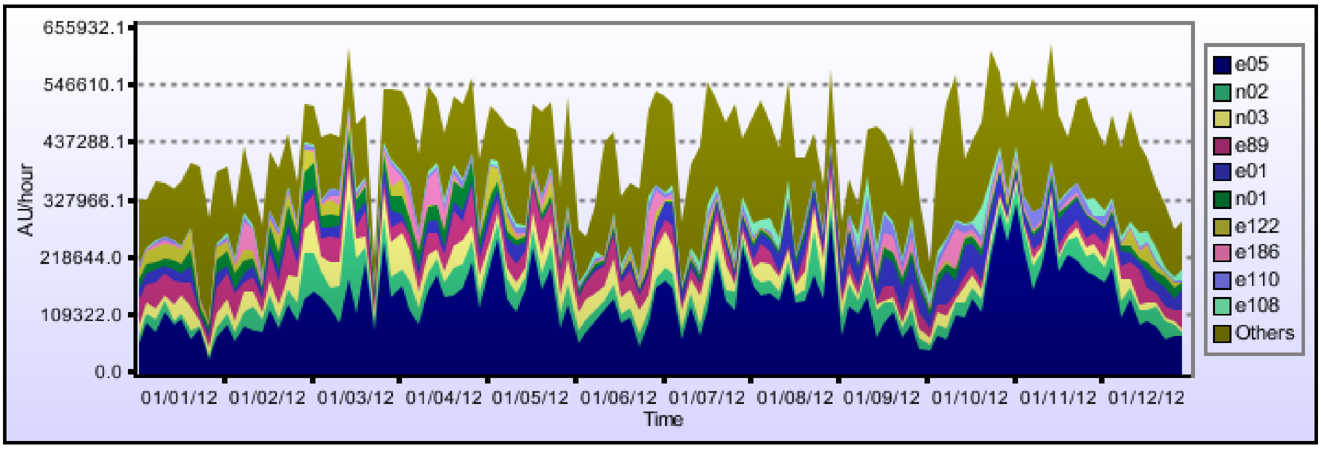

A total of 5,665,747.8 kAUs were available during this period. The utilisation was as follows:

Project	kAUs	Raw kAUs	Number of Jobs	Utilisation
y01	0.0	10.2	70	0.0%
y02	1310.8	31851.6	444	0.6%
y03	1603.2	1962.2	3,871	0.0%
y04	0.0	0.0	2	0.0%
y05	0.0	0.0	1	0.0%
y07	22.6	42.4	532	0.0%
z01	4671.8	8191.1	2,229	0.1%
z02	104.7	238.2	23	0.0%
z03	6010.4	11251.6	3,861	0.2%
z06	2.4	4.1	16	0.0%
z12	16.9	20.4	10	0.0%
Internal Total	**13742.7**	**53571.9**	**11,059**	**1.0%**
c01	15917.4	36798.1	8,196	0.7%
e01	86178.2	204282.0	12,638	3.6%
e05	658332.8	1215497.9	127,397	21.5%
e10	13498.9	25118.1	8,488	0.4%
e19	0.0	0.0	1	0.0%
e24	128.6	276.3	32	0.0%
e68	96.4	193.1	18	0.0%
e71	8757.0	14586.8	1,217	0.3%
e76	6405.8	10998.4	603	0.2%
e82	3286.9	5632.6	188	0.1%
e85	2521.9	3699.5	49	0.1%
e89	111401.2	245037.1	30,883	4.3%
e104	4907.2	8717.2	2,393	0.2%
e106	3430.8	6314.4	561	0.1%
e107	1062.9	7291.8	1,202	0.1%
e108	39748.9	70724.8	1,913	1.3%
e110	43062.3	80025.9	3,943	1.4%
e122	47901.6	85574.4	6,149	1.5%

Project	kAUs	Raw kAUs	Number of Jobs	Utilisation
e124	15276.6	56205.8	3,424	1.0%
e125	25539.5	43925.0	1,175	0.8%
e126	9249.2	15728.6	319	0.3%
e127	2986.2	5245.4	265	0.1%
e128	3197.7	5736.0	310	0.1%
e129	1930.4	3245.7	284	0.1%
e130	8547.7	14629.6	641	0.3%
e139	170.7	214.4	119	0.0%
e141	3992.5	7215.5	1,665	0.1%
e145	131.2	228.0	151	0.0%
e149	4847.6	10553.3	359	0.2%
e155	128.5	219.9	62	0.0%
e156	274.9	500.8	293	0.0%
e158	370.5	900.3	215	0.0%
e159	23.9	41.7	471	0.0%
e160	15734.4	29081.7	1,035	0.5%
e163	3738.9	6644.3	1,533	0.1%
e171	0.0	0.0	1	0.0%
e173	1381.0	2556.8	1,058	0.1%
e174	3459.7	6420.6	207	0.1%
e175	6429.7	42478.1	1,042	0.8%
e177	2824.9	5151.7	1,447	0.1%
e179	2364.9	3982.0	342	0.1%
e183	7323.3	12536.7	478	0.2%
e184	1213.4	2800.9	378	0.1%
e185	198.2	289.7	61	0.0%
e186	43220.0	80927.0	4,885	1.4%
e187	643.1	1044.9	1,392	0.0%
e192	3956.1	6782.5	370	0.1%
e193	13735.7	25039.6	1,263	0.4%
e202	14447.5	24733.7	493	0.4%
e203	4151.0	7324.6	962	0.1%
e204	1633.6	2852.2	126	0.1%
e206	0.5	0.8	51	0.0%
e207	18040.4	56834.1	120	1.0%
e210	33.6	62.7	70	0.0%
e212	0.0	0.0	2	0.0%
e213	9.7	16.7	81	0.0%
e214	2679.8	3351.0	310	0.1%
e215	5250.0	8241.8	499	0.2%
e216	1186.4	3038.5	991	0.1%
e217	1823.2	3013.7	784	0.1%
e218	310.7	394.4	139	0.0%
e219	0.0	58.0	1	0.0%
e220	17021.6	49873.2	1,072	0.9%
e224	419.0	673.4	202	0.0%
e225	220.6	411.9	82	0.0%
e228	547.3	968.3	271	0.0%
e229	10017.7	18539.1	6,889	0.3%
e230	932.7	1675.7	360	0.0%
e231	780.8	1500.4	527	0.0%
e232	290.7	521.0	239	0.0%
e233	7821.3	10225.3	1,272	0.2%
e234	1969.7	3463.5	526	0.1%

Project	kAUs	Raw kAUs	Number of Jobs	Utilisation
e235	455.2	830.4	630	0.0%
e236	613.3	1349.1	886	0.0%
e237	399.0	745.8	629	0.0%
e238	491.6	908.0	296	0.0%
e239	0.0	0.3	41	0.0%
e240	314.1	573.9	568	0.0%
e241	1.4	2.5	50	0.0%
e243	0.0	0.0	10	0.0%
e244	8.4	14.9	734	0.0%
e245	812.8	1479.6	533	0.0%
e246	0.0	0.0	1	0.0%
e247	220.9	412.9	77	0.0%
e248	102.4	206.7	54	0.0%
e249	2938.1	5294.1	382	0.1%
e250	1635.9	2899.7	49	0.1%
e251	11077.4	20359.5	256	0.4%
e252	7697.1	14100.8	742	0.3%
e253	14994.6	26859.9	433	0.5%
e254	60.0	205.2	2	0.0%
e256	37.9	64.8	72	0.0%
e257	1160.4	1986.5	116	0.0%
e258	3280.7	5631.6	107	0.1%
e259	9634.2	20401.2	120	0.4%
e260	1910.3	3270.2	1,194	0.1%
e261	0.6	1.1	40	0.0%
e262	1.2	2.1	63	0.0%
e263	2116.5	3623.1	218	0.1%
e264	117.9	201.8	14	0.0%
e267	0.1	0.2	28	0.0%
j01	8151.4	18102.2	2,138	0.3%
EPSRC Total	**1387350.4**	**2738397.0**	**257,668**	**48.3%**
n01	76150.1	130009.4	16,151	2.3%
n02	150882.7	268695.0	124,891	4.7%
n03	139754.6	246678.7	25,419	4.4%
n04	33722.1	59578.9	28,991	1.1%
u07	0.0	0.0	1	0.0%
NERC Total	**400509.5**	**704962.0**	**195,453**	**12.4%**
b09	0.5	0.9	21	0.0%
b14	19681.0	33985.8	2,634	0.6%
b100	8.3	15.5	152	0.0%
BBSRC Total	**19689.9**	**34002.2**	**2,807**	**0.6%**
p01	6878.9	11927.1	487	0.2%
STFC Total	**6878.9**	**11927.1**	**487**	**0.2%**
e168	79.5	238.8	148	0.0%
x01	27419.6	49831.2	17,716	0.9%
External Total	**27499.1**	**50070.0**	**17,864**	**0.9%**
b10	0.0	0.0	12	0.0%
d03	0.0	0.0	3	0.0%
d04	1.4	2.4	33	0.0%
d11	4245.0	7644.6	847	0.1%
d15	12.5	21.2	305	0.0%
d16	386.3	589.2	14	0.0%
d23	9.1	14.1	270	0.0%
d25	9622.6	17134.0	2,833	0.3%

Project	kAUs	Raw kAUs	Number of Jobs	Utilisation
d26	494.2	609.6	1,612	0.0%
d27	86.0	145.0	781	0.0%
d28	162.5	2014.7	858	0.0%
d29	571.4	1269.5	634	0.0%
d30	38.8	48.6	35	0.0%
d32	827.3	1541.6	1,622	0.0%
d34	171.8	298.9	2,425	0.0%
d37	8178.3	23734.2	3,690	0.4%
d38	410.1	744.4	515	0.0%
d39	986.1	1403.4	808	0.0%
d40	216.5	398.3	153	0.0%
d41	38712.9	63620.0	13,011	1.1%
d42	734.2	1372.3	87	0.0%
d43	1141.4	1954.0	286	0.0%
d45	29.6	36.7	1,279	0.0%
d46	0.2	0.4	8	0.0%
d47	0.1	0.2	3	0.0%
gd11	0.0	0.0	39	0.0%
i04	44.0	75.4	7	0.0%
x07	7.1	13.0	46	0.0%
Directors Time Total	**67089.7**	**124685.5**	**32,216**	**2.2%**
pr1u0702	1355.5	2320.5	102	0.0%
pr1u0704	8498.6	17185.4	1,011	0.3%
pr1u0705	14042.1	25429.3	815	0.5%
pr1u0706	14410.8	25959.7	431	0.5%
pr1u0804	15062.6	27225.6	4,882	0.5%
pr1u0805	9060.7	15510.9	422	0.3%
pr1u0806	370.7	634.7	103	0.0%
pr1u0807	75.4	129.0	95	0.0%
pr1u0808	13511.2	29845.8	54	0.5%
pr1u0809	1161.9	2355.4	38	0.0%
pr1u0902	532.1	910.9	130	0.0%
pr1u0903	12.0	20.6	9	0.0%
pr1u0904	133.6	228.7	44	0.0%
pr1u0905	51.1	87.4	74	0.0%
pr1u0906	1.4	2.4	12	0.0%
PRACE Total	**78279.6**	**147846.3**	**8,222**	**2.6%**
i03	3.4	5.8	31	0.0%
x08	35.8	66.9	45	0.0%
x10	19.4	36.4	83	0.0%
x11	3.4	6.2	14	0.0%
Industrial Total	62.0	115.3	173	0.0%
Total	**2001101.8**	**3865577.3**	**525,949**	**68.2%**

3.2.4 HECToR Utilisation by Application Area

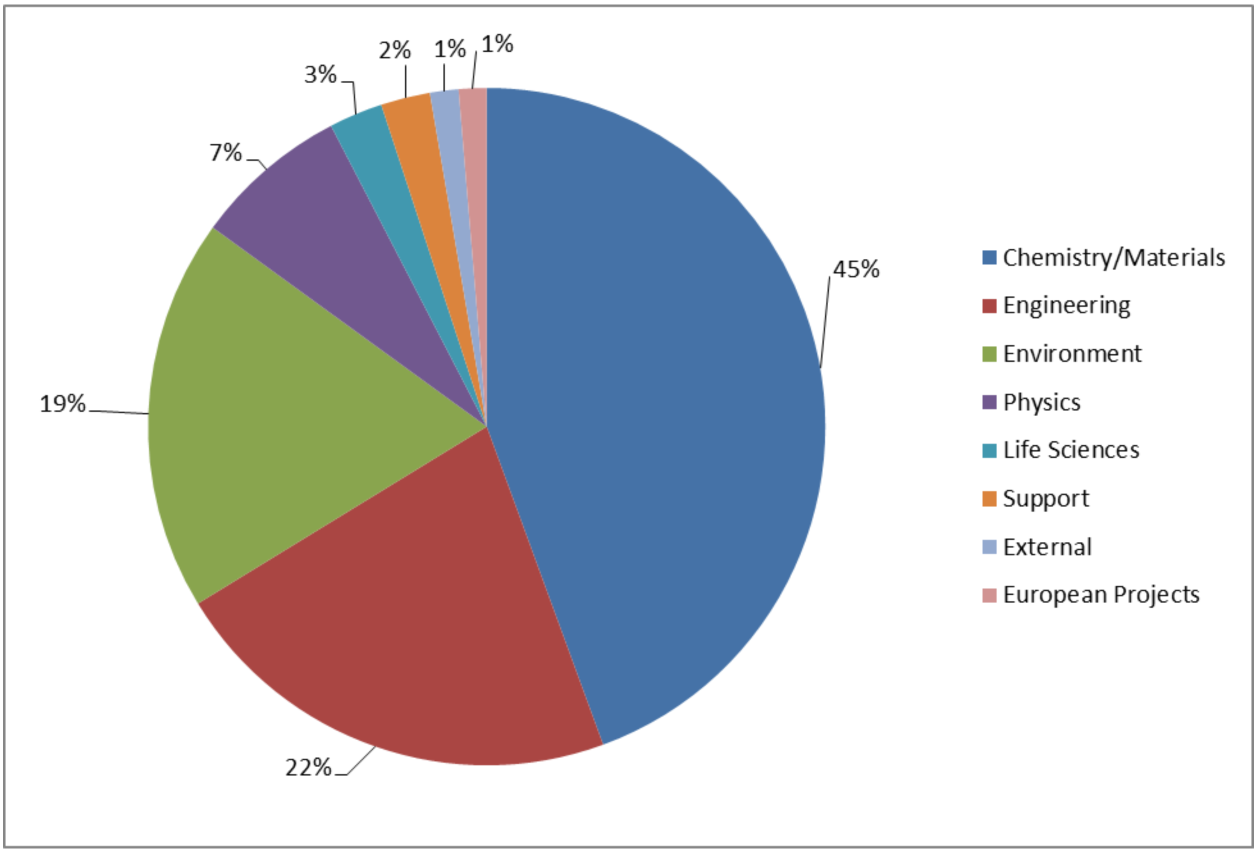

3.2.5 HECToR Code Usage Statistics

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

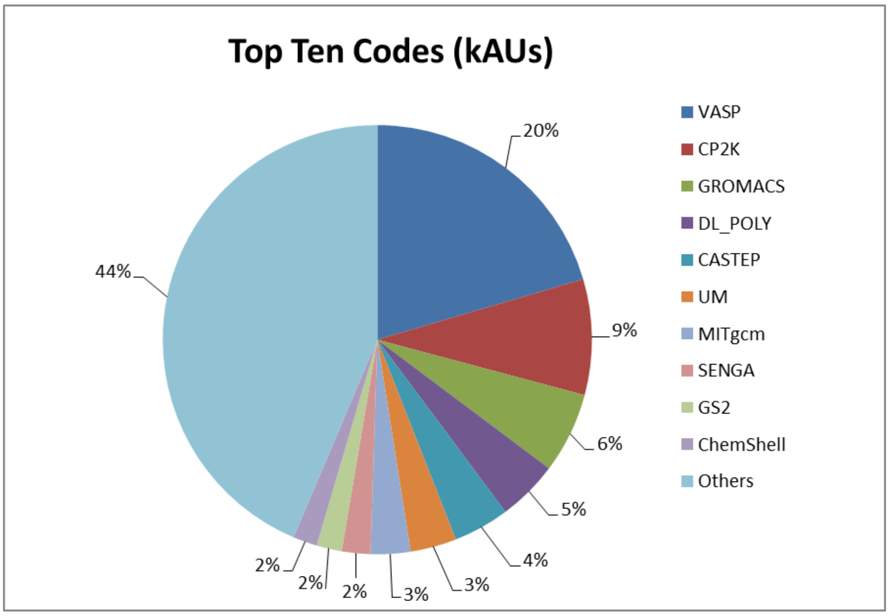

Code	% kAUs	%Jobs
VASP	20.4%	21.1%
CP2K	8.7%	6.0%
GROMACS	6.1%	4.0%
DL_POLY	4.6%	1.4%
CASTEP	4.2%	7.0%
UM	3.5%	6.8%
MITgcm	3.0%	15.8%
SENGA	2.2%	0.1%
GS2	1.9%	0.7%
ChemShell	1.8%	1.3%
NEMO	1.6%	0.3%
LAMMPS	1.5%	0.8%
ONETEP	1.0%	2.0%
CASINO	0.8%	0.1%
WRF	0.7%	0.6%

Code	% kAUs	%Jobs
NAMD	0.6%	0.6%
Fluidity	0.6%	0.3%
NWChem	0.6%	0.9%
CPMD	0.6%	0.4%
Terra	0.6%	0.2%
Quantum Espresso	0.5%	1.1%
Shelf	0.4%	2.2%
HELIUM	0.3%	0.0%
GLOMAP	0.2%	1.9%
GAMESS-US	0.2%	0.1%
Amber (PMEMD)	0.1%	0.2%
SIESTA	0.1%	0.1%
GAMESS-UK	0.0%	0.0%
MOLPRO	0.0%	0.0%
Others	33.2%	23.8%

3.2.6 HECToR Projected Utilisation

In relation to projected utilisation, analysis was completed in 2012 using the historic data in SAFE. This analysis looked at the AU spend profile of Class1a projects in relation to the project duration. Both the charged AUs and Raw AUs (before any discount) were graphed.

As below, the median shows that projects are approximately 25% through their expected duration before they start to spend AUs. Encouraging users to start quickly is critical either via pump–priming prior to Class1, or with CSE assistance.

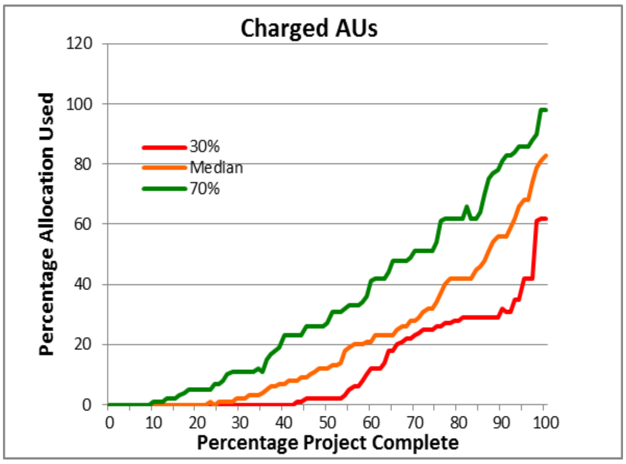

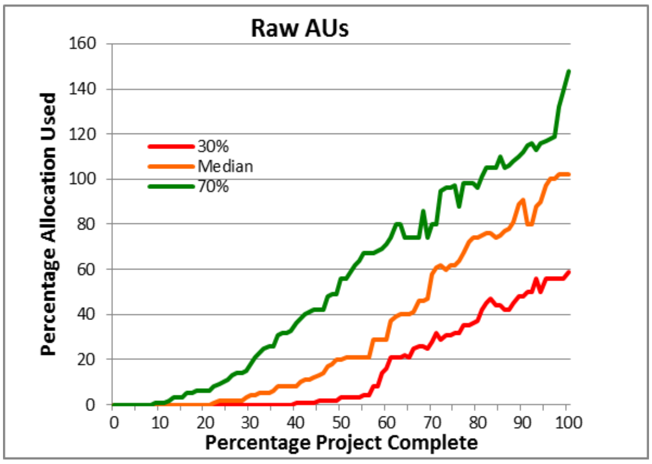

3.3 HECToR GPU

The HECToR GPU test system came online in March 2011. The initial allocation was shared 50% EPSRC and 50% Edinburgh. In total there have been just 4 EPSRC/NERC projects that have been setup on the GPU as part of the Class1b access mechanism. In addition a number of Edinburgh based projects have trialed codes. The purpose of the system is to investigate using GPUs as opposed to running major production runs.

Unfortunately due to a disk hardware failure the GPU was unavailable during November and December. The system has since been completely rebuilt, and work is ongoing to re-install the batch submission system.

From January to November the system was 38% utilised. In total, over 4600 jobs were run on the system.

Users	% Jobs	% Usage
Edinburgh	46%	73%
EPSRC	54%	27%

Similar to 2011, whilst the split in the volume of jobs has been fairly even between Edinburgh and EPSRC, the GPU usage has been dominated by one EPSRC project that is using a code already ported for the GPU. The majority of other jobs are very small test jobs.

3.4 Helpdesk

A total of 4648 queries with a specified service metric were completed in this period. There were also 886 queries with no metric completed in the same period. These 'no-metric' queries relate to non-service aspects such as the GPU, Nerc LMS and RDF. The volume of tickets raised regarding the LMS is very low (12 in 2012). On the RDF and GPU, the majority of requests have been for new user setups and quota configuration.

3.4.1 Helpdesk Targets

Metric	Pass	Total	Fraction	Target
All queries finished in 1 day	3829	3877	98.8%	97.0%
Admin queries finished in 1 day	3578	3621	98.8%	97.0%
Queries assigned in 30 min	4621	4648	99.4%	97.0%
Technical assessments in 10 days	129	130	99.2%	97.0%

3.4.2 Queries by Service Metric

Service Metric	Queries	Percentage
Automatic	2419	43.7%
Admin	1202	21.7%
In-depth	641	11.6%
Technical	256	4.6%
Technical Assessment	130	2.4%
No Metric	204	3.7%
Automatic Non-Service *	682	12.3%

**RDF, LMS, GPU User Admin Tickets*

3.4.3 Queries by Category

Query Category	Queries	Percentage
New User	641	13.8%
New Password	520	11.2%
Set user quotas	376	8.1%
Set group quotas	324	7.0%
3rd Party Software	314	6.8%
Access to HECToR	311	6.7%
None	240	5.2%
Disk, tapes, resources	231	5.0%
Batch system and queues	201	4.3%
Compilers and system software	199	4.3%
User behaviour	171	3.7%
Join Project	129	2.8%
User programs	124	2.7%
Login, passwords and ssh	121	2.6%
New Group	109	2.3%
Make Reservation	103	2.2%
Add to group	84	1.8%

Query Category	Queries	Percentage
Other	83	1.8%
Courses	53	1.1%
Create certificate	48	1.0%
SAFE	43	0.9%
Node Failure	43	0.9%
Update account	32	0.7%
Delete from project	31	0.7%
Archive	24	0.5%
Static website	23	0.5%
Performance and scaling	19	0.4%
Grid	15	0.3%
Remove account	14	0.3%
gpu	7	0.2%
Network	5	0.1%
Delete Certificate	5	0.1%
Delete from group	3	0.1%
Porting	2	0.0%

3.4.4 Queries by Handler Category

Handler	Total	Admin	Technical	In-depth	TA	Automatic	Automatic non-service	No Metric	Percentage
USL	1503	1040	176	172				115	27.2%
CSE	451	4	2	312	130			3	8.2%
OSG	3372	150	41	22		2419	682	58	60.9%
Cray	208	8	37	135				28	3.8%

3.5 User Quality Tokens

35 positive and 4 negative tokens were received in 2012. Corrective actions were taken to address the comments raised on the negative tokens.

3.6 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%	95.1%	100.0%	92.1%	100.0%	96.8%	99.5%	99.6%	100.0%	100.0%	100.0%	99.6%	99.6%	98.5%
Technology MTBF (hours)	100	126.4	244	∞	366	∞	732	366	366	∞	∞	∞	732	732	732
Technology Throughput, hours/year	7000	8367	7651	8648	7945	8681	8365	8603	8638	8654	8646	8393	8599	8693	8483
Capability jobs completion rate	70%	90%	93.3%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	95.5%	100.0%	100.0%	100.0%	100.0%
Non in-depth queries resolved within 1 day (%)	85%	97%	99.4%	98.9%	98.9%	100.0%	99.1%	99.2%	99.3%	97.3%	95.0%	98.8%	99.8%	98.4%	98.7%
Number of SP FTEs	7.3	8.0	8.8	8.6	8.9	8.4	8.8	8.2	7.8	9.1	8.4	8.5	8.7	7.3	8.5
SP Serviceability (%)	80.0%	99.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%

Colour coding:

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

4 Systems Hardware and Software

4.1 HECToR Technology Changes

In terms of the volume and impact of technology change – 2012 was a more stable year than previous years. Service disruption was kept to a minimum whenever possible.

4.1.1 Phase 3 Acceptance Testing

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 final 10 cabinets were added in January 2012, with acceptance testing and availability trials completed successfully thereafter.

4.1.2 RDF Integration

Long term data storage which outlives an individual national service is a key requirement for researchers. The UK Research Data Facility (RDF) was brought online in 2012 to fulfill this need. The RDF comprises of 7.8PB of disk storage, with 19.5PB of backup tape .The HECToR service was integrated in May, with key consortia from the EPSRC and NERC communities being granted storage allocations on the service.

4.1.3 Operating System Upgrade

The HECToR Operating system was upgraded to CLE4.0 (update 3) in June. Multiple patch sets were also installed to provide fixes and improved functionality throughout the year

4.2 Service Improvements

4.2.1 24 Hour Queues

In response to user requests for longer jobs, 24 hour queues were implemented on a trial basis in July 2012. 24 Hours jobs accounted for 16% of the utilisation from July to January. The initial 6 month trial completed successfully in January 2013, and the decision was taken by the HECToR Scientific Advisory Committee (HSAC) to adopt this as part of the service.

4.2.2 DDT Debugger

The DDT Debugger was installed on HECToR in 4Q12. The software license was funded by UoE HPCx and will benefit all users. A number of DDT training courses are planned in 2013.

4.2.3 Status Reporting Pages

Additional functionality was added to the system status page on the HECToR website in 2012. This provides users with a high level overview of the status of the job queues.

4.3 Severity-1 Incidents

4.3.1 Technology Failures

Cray supplied technology has been responsible for 12 severity-1 incidents in 2012. Eight of these incidents occurred in the first half of the year with four in the second half of the year.

This is a breakdown of the failure categories:

- Four instances of late return to service following scheduled maintenance sessions
- Three PBS batch subsystem related failures
- One cabinet blower failure
- One lustre filesystem failure affecting EPSRC users
- One instance of HSN failure
- One HSN failure during a ‘warmswap’ maintenance operation
- One incident due to Phase 3 acceptance testing overrun.

Detailed and careful planning is carried out ahead of scheduled maintenance sessions, unfortunately a late return to service following these sessions have occurred on four occasions in 2012. We take the opportunity to analyse the circumstances surrounding these incidents and any lessons learned are built into site procedures to minimise future risks.

Of the three PBS related failures, two were due to PBS server crashes which are still under investigation by Altair. The third issue occurred when a user process consumed excessive memory on the job launch nodes. Configuration changes have been made to mitigate this type of failure in future.

The temporary loss of lustre occurred following an error on the management server for the file system. This failure did not impact all users of the system. Memory components suspected of triggering this problem have been removed from the server configuration.

An attempt to use an online maintenance feature ‘warmswap’ to repair a module failure just prior to running the Phase 3 acceptance tests ran into an end-case failure which resulted in a non-routable network. New software was installed in April that fixes the problem encountered.

An overrun of the allocated time for Phase 3 acceptance testing occurred due to a software bug in the HPL application.

4.3.2 Service Provision Failures

There were no failures attributed to the service provision in 2012.

4.3.3 External Failures

There were no external failures in 2012.

4.4 Single Node Failures

The number of single node fails causes no concern. The peak of node fails in October was related to a software bug.

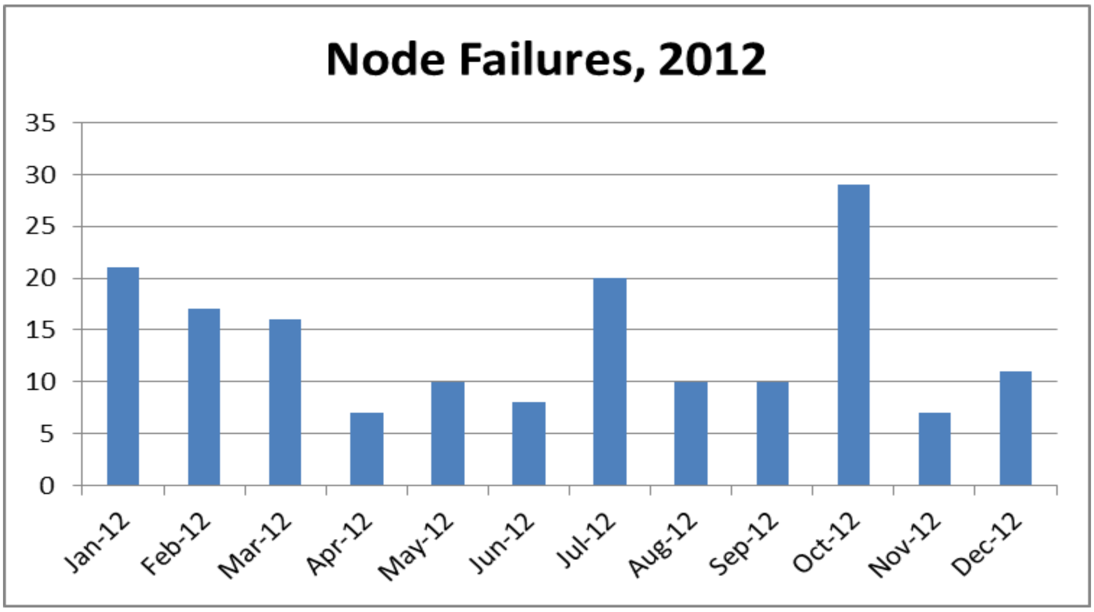

5 Cray Centre of Excellence

5.1 Executive Summary

The beginning of the year saw the acceptance of Phase 3 of the HECToR service, an upgrade and expansion of the existing Cray XE6 system from AMD's Magny-Cours processors to the latest generation of AMD Interlagos chips. This represented a significant redesign of the underlying processor architecture that provided a significant increase in the total available computing power for HECToR's users.

In support of the upgrade the Cray Centre of Excellence (CoE) and Exascale Research Initiative (ERI) have engaged in user training, workshops and high level support of users. They have also been involved in promoting the HECToR service throughout the UK, Europe and the rest of the world. Members have contributed to the dCSE process, provided guidance and support for students participating in High Performance Computing training. The CoE staff have continued to build and maintain close relationships between Cray's Application and Development teams and coordinate with HECToR activities where specialist input was essential.

5.2 Applications Support

Each of Cray's HECToR teams, the CoE, the Cray service team and the ERI have been actively preparing HECToR's service and user support teams for the upgrade, providing targeted insight and advice on how to get the best from this newest generation of supercomputer. The introduction of the system into full production heralded the start of an intensive period of support as users sought to fully exploit the latest incarnation of the HECToR service.

As well as providing support to HECToR users, the CoE team maintains close links with the network of other Cray Centres of Excellence at leading supercomputer sites across the Americas, Europe and Asia sharing best practice in user support and domain experience between sites. The Cray Exascale Research Initiative provides close links with Cray's Research and Development teams and allows HECToR's users and support teams to contribute to the discussion of next generation HPC software and hardware.

5.2.1 UPSCALE

Through the international network CoEs, the HECToR team were invited to assist scientists from the National Centre for Atmospheric Science (NCAS) at the University of Reading and the Met Office, already experienced HECToR users, in porting and optimising the Unified Model on HERMIT (another Cray XE6 system installed at HLRS, Stuttgart).
The “UK on PRACE: weather-resolving Simulation of Climate for globAL Environmental risk” (UPSCALE) project received a large allocation of computer resources on HERMIT, a PRACE Tier-0 system, to run high resolution investigations into the potential impact and frequency of extreme weather events. This required extending a suite of 10 models configured at forecast resolution, typically only used to simulate a few days or a week, to run for over 25 model years. This presented new challenges in scaling the application and managing the large volume of data generated.

The Centre of Excellence team embarked on a three month project to optimise the model at for large scale runs which resulted in a 39% improvement in runtime performance versus the initial port. Areas optimised covered task placement, cache optimisations and compiler optimisation. It was felt the CoE team was best placed to assist due to the combined scientific and technical domain expertise available.

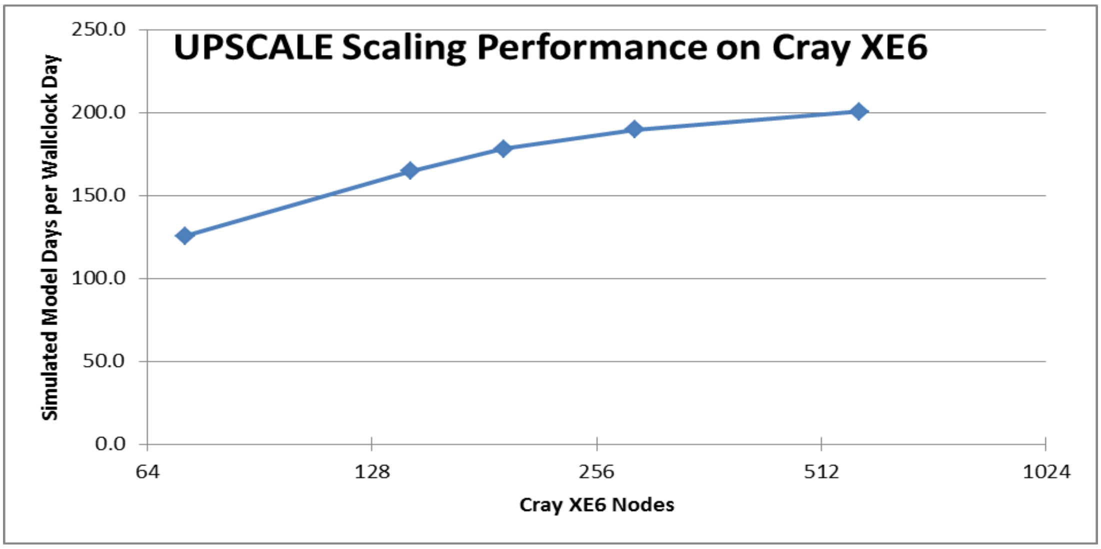

5.2.2 Ludwig

As part of Cray and EPCC's on-going collaborative research commitment to preparing applications for exascale the Ludwig code was ported to the Cray XK6 system at Oak Ridge National Laboratory. Ludwig is a lattice Boltzmann code which is part of HECToR's core capability which is used for understanding and controlling phase separations in liquid mixtures. A practical example is use in understanding the effect of ice crystals on the shelf life of frozen desserts.

Members of the HECToR team successfully ran the application over 10 cabinets of the Cray XK6 using 936 GPUs and simulating a lattice of 2000^3 grid points. The deviation from linear scaling was less than 5% per cent at the highest scales and overall the application performed two-and-a-half times faster than a comparable x86 system.
Details from the run were presented at NVIDIAs GPU Technology Conference by EPCC and featured on an NBC news broadcast in the United States.

Ludwig Scaling Graph

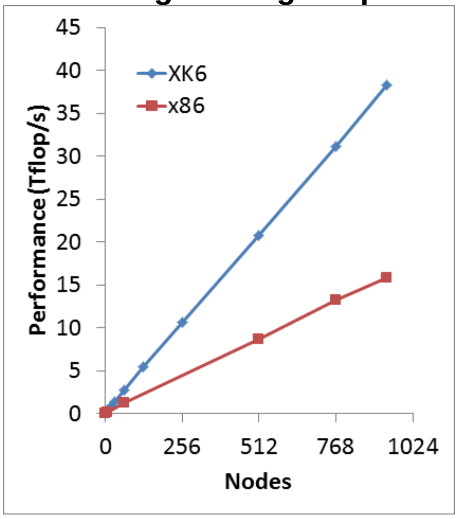

5.3 Seminars

Members of Cray's teams at HECToR have attended a variety of seminars and meetings through 2012, representing HECToR within the international HPC community.

- The CoE delivered a paper titled “Applying Automated Optimisation Techniques to HPC Application” at the Cray User Group meeting at HLRS in Stuttgart in May 2012.
- Cray's keynote address at the Isaac Newton Institute’s “Weather and Climate Prediction on Next Generation Supercomputers” workshop satellite meeting held at the Met Office in October 2012 was given by a member of the CoE.
- Members of the CoE represented Cray at the dCSE Review Panel in June 2012 and the Support Technical Meeting in September 2012 as well as the HECToR town hall meeting in December.
- The ERI were invited to provide a Cray tutorial at Supercomputing 12 on OpenACC, an open standard for Accelerator based programming.
- Cray provided a decommissioned supercomputer blade for display and assisted EPCC's at “Supercomputers and You” exhibition stand at the British Science Festival 2012 in Aberdeen.

5.4 Workshops and Training

As well as representing Cray and HECToR around the world, HECToR's Cray teams have provided numerous training opportunities for academics interested in new and emerging technologies. These include:

- Providing numerous workshops on optimisation for CPUs and GPUs across Europe:
    - Two HECToR specific two-day Cray CoE XE6 Optimisation Workshops in Edinburgh (July), and for the first time, Reading (November). These were run in collaboration with the PRACE Advanced Training Programme and enabled users to engage directly with experts from Cray and Roguewave, suppliers of the Totalview debugging suite.
    - Also as part of the PRACE scheme, Cray assisted in organising, in collaboration with EPCC, a two day workshop in August for GPU programmers.
    - Users from the University of Manchester requested and were provided with a specific one day training course on OpenACC and GPUs.
    - Team members have been involved with general training for HPC users at sites across Europe:
        - CSCS (Switzerland), PDC/KTH (Sweden) and HLRS (Germany).
- Members were involved, co-supervising one of EPCC's MSc students. The teams also provided additional lecture and practical materials on PGAS programming methods.

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 more than 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 to HECToR, 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 2012

2012 started with the launch of Phase 3 of the service, which went relatively smoothly. All our documentation and training material was revised to reflect the new system. Over the year the Core CSE team has continued to assist users, closing 550 queries in the period. The CSE training programme remains popular and this year for the first time we delivered courses at the Universities of Liverpool and Plymouth. To date we have filled over 1600 places on our training courses, and we continue to receive excellent feedback from attendees and invitations to return to institutions to provide further training in future. EPSRC has provided additional DCSE funding allowing us to continue the programme until September 2013 (the original service end date). In addition further funding has been provided to offer DCSE to non-HECToR users in a pilot programme which started in the summer. While some of the eight funded projects cover similar topics to conventional DCSEs (e.g. introduction of hybrid parallelism and improvements to scalability), others include adapting applications to make use of GPUs, using MATLAB® as a front-end to a parallel application, and agent-based modelling for synthetic biology. To date the DCSE programme has allocated over 65 person years of effort to improving codes deemed important to UK research. We continue to feed code improvements back to code owners so that the worldwide HPC community can benefit.

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, they will be re-assigned to the Service Provider (UoE HPCx) or to Cray. The following chart shows how many queries were opened and closed each month.

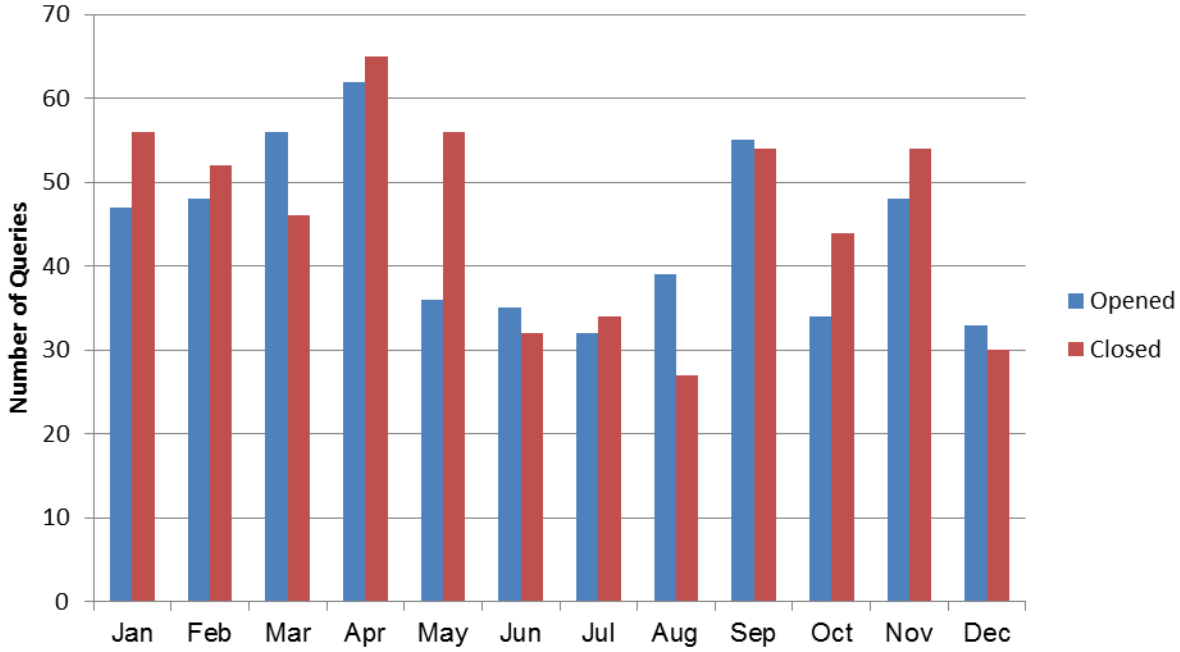

In cases where a technical query (as opposed to a request for a Technical Assessment) is resolved by the team (strictly speaking where the query is closed within the SAFE system by the CSE team), the user is invited to fill in a questionnaire giving feedback about his or her experience and satisfaction with the outcome. This year, 99 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:

Q1: Has the problem in your query been resolved by the information proved by the Helpdesk?

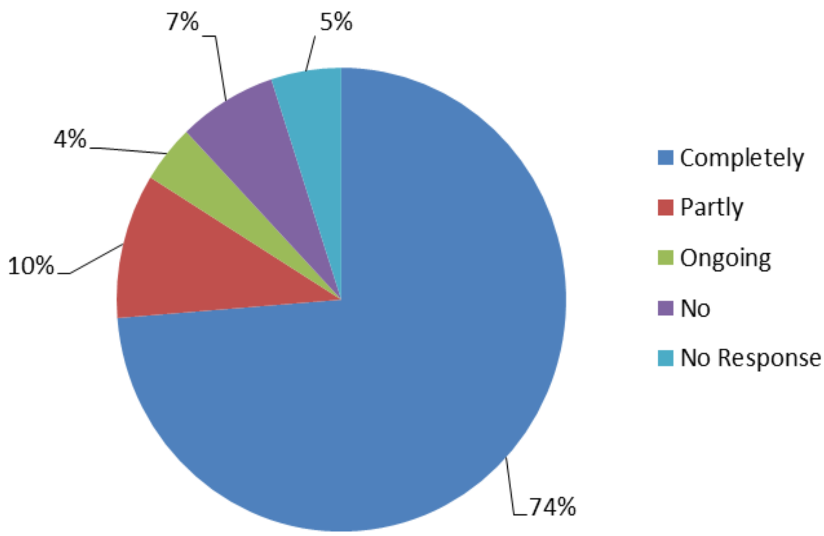

Q2: Were all communications replied to promptly and the text clear and understandable?

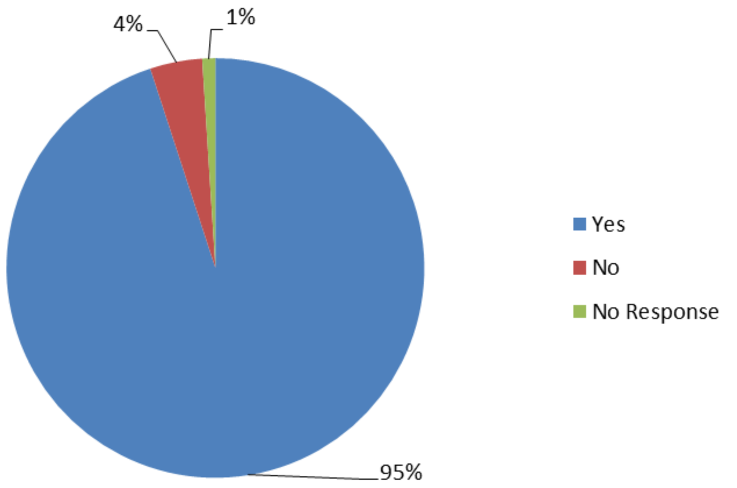

Q3: Were technical issues explained in sufficient detail?

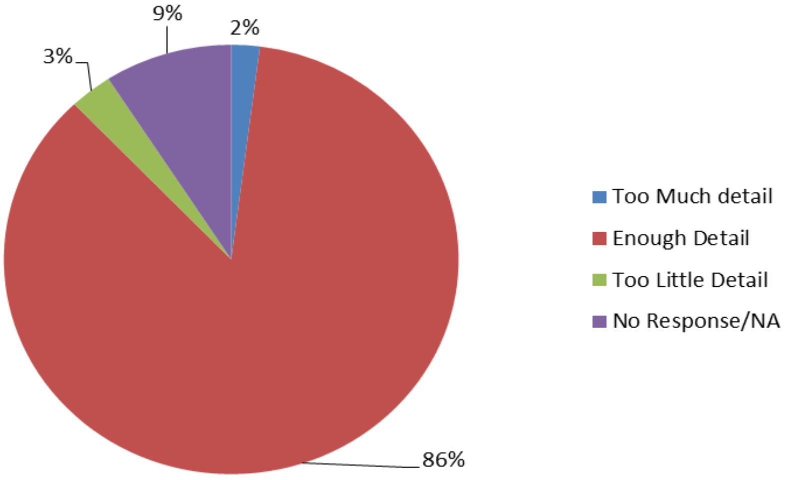

Q4: Did you use the HECToR CSE documentation to try and find a solution before submitting a query to the helpdesk?

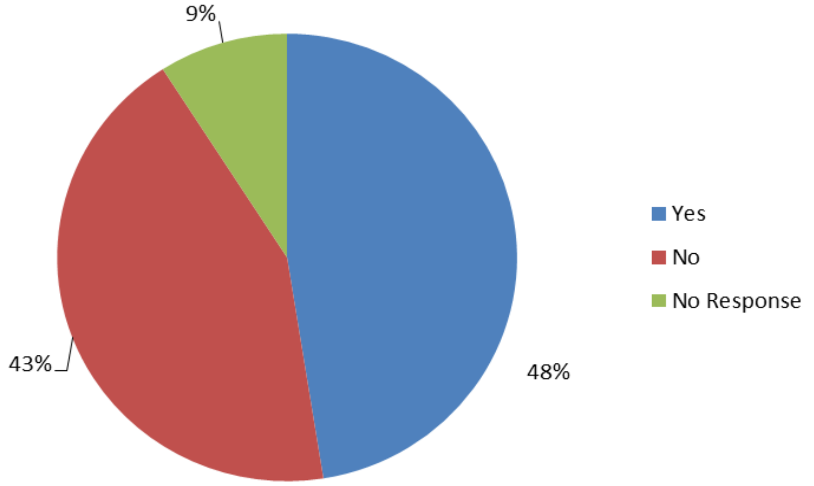

6.4 Training

During this period NAG has offered a range of training courses, most of which were not HECToR-specific. Most courses have been run by NAG staff, although we have delivered courses in collaboration with the Universities of Warwick, Bristol and Oxford, and with STFC. 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 Culham, Imperial College, Bath, Leicester, Sheffield, Plymouth, Warwick, UCL, Southampton and Liverpool, as well as at the NAG offices in Oxford and Manchester. A total of 72 days of training have been offered and delivered to 365 attendees.

The complete list of courses offered in 2012 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.

- An Introduction to CUDA Programming
- An Introduction to OpenCL Programming
- Core Algorithms for High Performance Scientific Computing
- Debugging, Profiling and Optimising
- DL_POLY
- Fortran 95
- Introduction to HECToR
- Multicore
- OpenMP
- Parallel Programming with MPI
- Parallel IO
- Scientific Visualisation

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 HECToR Distributed CSE

The DCSE panel met three times during 2012, and considered a total of twenty-four applications requesting 174 months of effort. Of those, nineteen were recommended to receive a total of 133 months of effort.

Nineteen projects ended during 2012. At the end of December fourteen projects were running with six more scheduled to start in 2013. Since the start of the service the DCSE panel has allocated over 65 person-years of effort. In September NAG organised a workshop for all DCSE participants in Manchester, slides from which are available on the HECToR website.

6.7 DCSE Projects Completed in 2012

6.7.1 Adding Spin Polarisation and van der Waals Energy Correction to Conquest

Conquest is a linear scaling code which is used for calculations in metallic systems. This second DCSE project for Conquest was allocated 12 months support at the June 2010 call, to extend functionality and enable scalable simulations of bio-molecular systems and magnetic systems, where van der Waals (vdW-DF functional) interactions are essential. Dr David Bowler of the Department of Physics and Astronomy at UCL supervised the work, performed by Lianheng Tong, also of UCL between February 2011 and February 2012.

The capability to perform both spin polarised calculations and van der Waals interactions was implemented. Results for the spin implementation were validated against known experimental values and calculations from the SIESTA *ab initio* molecular dynamics package. Furthermore, since existing Conquest kernels for distributing data and computation were used, the linear scalability was generally not affected. However, scalability of Conquest with the vdW-DF functional was found to be dependent upon the use of memory and FFTs within the code, since Conquest currently uses its own FFT implementation rather than an optimised FFT library. This work has been submitted to the Conquest code repository and is used within groups at UCL and Imperial College. On HECToR, around 5 million AUs per year are consumed by Conquest.

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

The R-matrix codes RAD, ANG and HAM are used to investigate electron atom collisions for the inner region. They build up (and diagonalise) the full Hamiltonian matrices for the collision problem, along with associated multipole matrices, in a finite volume. This project was allocated 12 months effort at the June 2010 round and followed on from a previous DCSE project which targeted the Outer Region R-matrix code, PRMAT. Dr Martin Plummer, Dr Andrew Sunderland and Professor Cliff Noble of STFC performed the work between April 2011 and March 2012. Dr Plummer also supervised the project.

The common objective to upgrade the RAD, ANG and HAM (PDG) codes for improved OpenMP and task management was achieved. RAD now scales to more than one node by using mixed-mode OpenMP with MPI. ANG (and HAM), were previously both serial codes; these now scale to hundreds of cores enabling collision problems to be tackled which were previously impossible. The MPI-2 parallelism with passive remote memory access takes care of load-balancing problems that occur in coarse-grained parallelism over loops with unpredictable amounts of work to perform. Furthermore, the use of FORTRAN 2003 features has improved the clarity of coding. The RAD, ANG and HAM codes are used within the UK-RAMP consortium which holds around 130 million AUs on HECToR.

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

VASP is used to perform *ab initio* electronic structure and quantum-mechanical molecular dynamics simulations on HECToR. This was the second DCSE project to improve the performance of VASP and was allocated 12 months support at the November 2010 call, the work was aimed at enhancing the hybrid density functional theory (DFT) performance of the code. Dr Scott Woodley of the Department of Chemistry at UCL supervised the work, which was performed by Andy Turner of EPCC between November 2010 and September 2011 and Gavin Pringle, also of EPCC, between October 2011 and October 2012.

OpenMP parallelisation has been added to the Projector-Augmented Wave routines and the hybrid DFT routines. This has enabled VASP to utilise all cores on a HECToR node more efficiently for runs using more than 8 nodes. Plans are in place to make these developments available to all HECToR users of VASP for their atomic scale materials simulations. VASP is currently the single most heavily used application on HECToR, consuming over 20% of the machine's overall resource.

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

CASTEP is a widely used DFT code. This fourth DCSE project for CASTEP was allocated 10 months effort at the June 2010 round and was supervised by Dr Keith Refson of STFC. The first objective was to both improve the MPI Buffer memory and optimise the I/O. The second objective was to further develop the band-parallel capability of CASTEP for multi-core architectures.

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, worked on the improvements to the band-parallel capability of CASTEP from October 2011 to January 2012. This work has improved the scalability of the code and the 2-way band-parallelism is now a viable option for CASTEP use on HECToR. These developments have also been introduced within the main CASTEP code base. CASTEP is currently the fourth most heavily used application on HECToR, consuming over 4% of the machine's overall resource.

6.7.5 CP2K - Sparse Linear Algebra on 1000s of cores

Following on from two previous DCSE projects for the CP2K atomistic and molecular simulation code, the aim of this work was to reduce the time to solution for typical simulations by implementing more efficient and scalable sparse matrix operations. The project was allocated 6 months support at the June 2010 round and was supervised by Dr Ben Slater of the Department of Chemistry at UCL. Iain Bethune of EPCC performed the work between January 2011 and January 2012.

The Distributed Block Compressed Sparse Row (DBCSR) matrix multiplication library, which is embedded within CP2K, was optimised. The 6144 atom test calculation of liquid water is now up to 20% faster using six threads and up to 2304 cores. A specialised small matrix multiplication library (libsmm) was also developed which employs an autotuning (loop reordering/unrolling) approach to produce specialised DGEMM routines, and can achieve a 10 x speedup for certain small matrices. Furthermore, the CSR index re-ordering routine was re-written, and a number of smaller OpenMP optimisations were implemented. For a typical matrix, this is now 8% faster for 128 MPI x 2 OMP, and 43% faster for 32 MPI x 8 OMP. All developments have gone into the main CP2K code base. CP2K is currently the second most heavily used application on HECToR, consuming over 10% of the machine's overall resource.

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

This was the third Fluidity (ICOM) DCSE project, with the overall aim to enhance the multi-core performance of the code, thus enabling increased scalability and efficiency on HECToR. The project was allocated 12 months effort at the June 2010 round. Dr Andrew Sunderland of STFC supervised the work and Xiaohu Guo, also of STFC, performed the work between October 2010 and September 2012.

Mixed-mode MPI/OpenMP parallelisation was implemented for the finite element assembly stage, linear preconditioners and solvers. This has decreased the overall memory footprint required per node, and also the amount of I/O operations required, enabling more efficient utilisation of HECToR. Plans are in place to make these developments available to all users of Fluidity. Over the last two years Fluidity has consumed nearly 6 million AUs on HECToR within Imperial College, which is in addition to over 10 million AUs used by the centrally installed version.

6.7.7 Developing NEMO for Large Multi-core Scalar Systems technology

NEMO is an ocean model which has been used successfully for many years in global and ocean basin simulations. Recently it has started to be used as a shelf-sea model. This was the second NEMO DCSE project, with the overall aim to improve performance by implementing techniques previously proven in POLCOMS. The project was allocated 12 months support at the June 2010 round and was supervised by Dr Stephen Pickles of STFC. The work began in January 2011 and completed in June 2012: Dr Pickles performed the work up to the end of December 2011, and then Dr Andrew Porter, also of STFC, took from January 2012 until June 2012.
Array index re-ordering and loop-nest optimisations were implemented so that the arrays now have the level index first. This has improved cache utilisation for the tri-diagonal solver, enabling a 3% increase in the single-core performance for the GYRE (no land) deep ocean test case, together with slightly improved scalability (due to the more efficient halo exchanges). Recursive k-section (multi-core aware) partitioning was also implemented for the domain decomposition. For the AMM12 test case, the new code performs slightly better than the original for up to 128 MPI processes. During the project a performance overhead of using REWIND with the Lustre file system on HECToR was identified and resolved, now giving a general 10% saving in NEMO production work. Developments from this project are in a branch of the NEMO source and will also be used to inform strategic decisions about future NEMO developments. For 2011-2012, HECToR usage of NEMO was around 25% of the total of NERC's allocation, which amounts to around 5% of HECToR’s overall resource.

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

UKRMol-in is part of a polyatomic suite of codes that are used to model electron- and positron- molecule scattering processes. These codes use *ab initio* methods for the solution of the time dependent Schrödinger equation and are based on R-matrix theory for electron scattering. The project was allocated 12 months effort at the December 2010 round and was supervised by Dr Jimena Gorfinkiel of the Department of Physics & Astronomy at the Open University. Michael Lysaght of the Open University worked on the project from April 2011 until October 2011. Paul Roberts of NAG took over the work in November 2011 and completed the project in May 2012.
In order to develop a scalable parallel version of UKRMol-in, the Hamiltonian build and diagonalisation steps were implemented with calls to SLEPc and PETSc. Further work was required to improve the load balancing and initial performance. Scalability to several nodes on HECToR is now achievable, allowing much larger problem sizes to be investigated. This work has been submitted to CCPForge and is used on HECToR within the UK-RAMP project.

6.7.9 CABARET on Jet Flap Noise and Quasigeostrophic Ocean Circulation Models

CABARET is used on HECToR both for modelling the aeroacoustics of complex jets, and for Quasigeostrophic (QG) ocean simulations. The CABARET algorithm constitutes a substantial upgrade of the second-order upwind leapfrog method and is very suited to HPC since it has a very local computational stencil. 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 supervised by Dr Sergey Karabasov of the School of Engineering and Materials Science at Queen Mary University of London. The project objectives were to improve the performance of the aeronautics version of CABARET and optimise a scalable implementation of the QG version. Phil Ridley of NAG performed the work between August 2011 and July 2012.

The parallel Helmholtz solver in the QG version of CABARET was optimised by implementing calls to hardware optimised global collectives. Initially, the code did not scale above two HECToR nodes for a representative 3 isopycnal layer problem with a 1025^2 grid. The code now shows good strong scalability to 256 cores and almost linear weak scaling for a fixed number of grid points on a per core basis. Furthermore, MPI-IO was implemented for the restart and output data in the aeronautics version of CABARET. This has reduced the amount of time required for post-processing and will enable more efficient utilisation of the file system on HECToR. Over the last two years, CABARET has consumed over 50 million AUs on HECToR, and a further 20 million AUs will be used for CABARET simulations during early 2013.

6.7.10 Scalable coupling of Molecular Dynamics and Direct Numerical Simulation of multi-scale flows Part 1

The Transflow CFD code and the Stream-MD molecular dynamics code both scale very well to over 800 cores on HECToR. The overall objective of this project was to implement a fully coupled and scalable version of both applications to enable the simulation of molecular dynamics with CFD. This project was allocated 6 months effort at the December 2010 round and the work was supervised by Dr Tamer Zaki of the Department of Mechanical Engineering at Imperial College. Lucian Anton of NAG performed the work between April 2011 and March 2012.

A coupled continuum MD model was developed for the TransFlow CFD solver and the Stream (MD) solver. The coupler includes a communications infrastructure (for internal data management) and interface module (for data access between applications). For a Lennard-Jones system of 3,317,760 molecules, scalability of the coupled application is good for up to 800 cores. The application is used on HECToR within Dr Zaki’s group under a resource allocation from the UK Turbulence Consortium.

6.7.11 Direct Numerical Simulations (DNS) of turbulent fluid flows

Incompact3D and related codes Compact3d and QuasiCompact3D are used to simulate complex fluid-flow problems by means of Direct Numerical Simulation. This DCSE project was a continuation of the first project which focused on improving the scalability of the FFTs in Incompact3D and was allocated 6 months support at the December 2010 round. The work was supervised by Professor John C. Vassilicos of the Department of Aeronautics at Imperial College. Ning Li of NAG performed the work between March 2011 and February 2012.

The objectives of the project were to port Compact3d and QuasiCompact3D to HECToR, and to implement a new communication library for efficient particle tracking simulations. Scalable versions of Compact3d and QuasiCompact3d have been produced which can be used to investigate the acoustic properties of multi-scale objects and also to perform high-fidelity simulations with particle-turbulence interaction effects. These applications regularly use 8,000-16,000 cores on HECToR for production runs, and so far have used over 51 million AUs.

6.7.12 Improved Data Distribution Routines for Gyrokinetic Plasma Simulation

The GS2 code is used to study low-frequency turbulence in magnetised plasma. Typically this is for assessing the microstability of laboratory produced plasmas, and to calculate key properties of the turbulence which results from instabilities. This project was allocated 6 months effort at the December 2010 round and is the second DCSE project concerning GS2, following on from the previous project to upgrade the FFTs. The project was supervised by Dr Colin Roach of the EURATOM/Culham Centre for Fusion Energy and the main objective was to improve the method of indirect addressing for data access. Adrian Jackson of EPCC performed the work between March 2011 and January 2012.

Performance of the local data copying associated with the data transform between the linear and non-linear calculations has been improved by around 40-50%, by replacing the costly indirect addressing with direct access mechanisms. For the non-linear calculations a new decomposition was implemented which allows varying amounts of data to be allocated to each process, and so reduces communication costs at large process counts. For a representative benchmark, a reduction in overall runtime of up to 17% can now be achieved. In particular, with 512 cores an overall 7% reduction in run time may be achieved and with 2048 cores this is nearly 20%. Around 150 million AUs have been allocated for use with GS2 and related applications on HECToR.

6.7.13 RMT for High Harmonic Generation

This was 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. RMT is a new *ab initio* method for solving the time-dependent Schrödinger Equation for multi-electron atomic systems in intense short laser pulses. The code was developed as part of the HELIUM DCSE project and the aim of this work was to add new functionality for the calculation of High Harmonic Generation (HHG) processes. The project was allocated 6 months of support at the June 2011 round and was supervised by Professor Ken Taylor of the Department of Applied Mathematics & Theoretical Physics at QUB. Jonathan Parker, also of QUB, performed the work between October 2011 and March 2012.

Tests to validate the calculation of the energy spectrum of HHG processes showed that the HELIUM and RMT codes gave results that agreed over a range of 9 orders of magnitude. The RMT code demonstrated speedup to about 1000 cores on HECToR for a small model of the helium atom. Furthermore, the two methods show good agreement for a series of HHG calculations in the frequency range 0.1168-0.94 atomic units. This project will enable the RMT code to utilise HECToR for a rigorous theoretical analysis of recent experimental advances.

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

SPRINT is an add-on package for the R statistical application. This was the third SPRINT DCSE project; previously the parallel Pearson correlation and permutation testing functions were implemented along with the randomForest decision tree classifier and rank product. The aim of this work was to implement bootstrapping and support vector machines (SVM) for use with SPRINT on HECToR. The project was allocated 6 months effort at the June 2011 round and was supervised by Terry Sloan of EPCC, with Michal Piotrowski of EPCC performing the work between February 2012 and July 2012.

The work has already enabled The Division of Pathway Medicine at Edinburgh to embark on a study of oscillatory behaviour and causality relationship of dynamic changes in expression profile in a time course transcriptomic experiment. This one experiment involves 3 biological conditions, with measurements every 30 minutes for 12 hours. The resulting data analysis involves 259 x 10^9 correlations and would not be possible without the new functionality in SPRINT and the capability of HECToR.

6.7.15 Improved Scaling for Direct Numerical Simulations of Turbulence

SWT is used for the direct numerical simulation of turbulent flow in an infinite plane channel and turbulent Couette-Poiseuille flow. SS3F solves the incompressible Navier-Stokes equations in the Boussinesq approximation using a 3D-Fourier representation. The aim of this project was to remove the limit on parallel scalability in both codes by improving their parallel FFT implementation. The project was allocated 6 months support at the June 2011 round and the work was supervised by Professor Gary Coleman of the Department of Engineering and the Environment at the University of Southampton. Roderick Johnstone, also of the University of Southampton, performed the work between November 2011 and July 2012.

For both codes, a 1-D cosine transform was implemented from FFTW3. For a 128×720×1440 mode problem, SS3F showed a 51% reduction in wall clock time per time step and a 55% reduction for 256×1440×2880 modes. For a 3072×325×1024 mode problem on 6144 cores, SWT showed a 53% reduction. Furthermore, a parallel 2-D-FFT/1-D-Chebyshev transform was implemented to work with the 1-D cosine transform, for both SS3F and SWT. For a 768×1536×3072 mode test problem, SS3F now scales to over 12000 cores with good efficiency. For a representative 3072×325×1024 mode problem, SWT now scales to 8192 cores with good efficiency. In general there is a 34% improvement for both codes, which are used under the UK Turbulence Consortium allocation on HECToR.

6.7.16 Improving Scaling in Conquest for HECToR Phase 3 - OpenMP-MPI Hybrid Implementation

The aim of this third Conquest DCSE project was to develop a new OpenMP / MPI version of the code to enable efficient scalability for up to 4,096 HECToR cores. The project was allocated 8 months effort at the June 2011 round, and was supervised by Dr David Bowler of the Department of Physics and Astronomy at UCL. The work was performed by Lianheng Tong, also of UCL, between February 2012 and September 2012.

OpenMP parallelisation has been successfully implemented in the matrix multiplication subroutines. The hybrid OpenMP/MPI implementation significantly improves the weak scaling of the code when compared with the original pure MPI implementation. Linear weak scaling is now achievable to several thousand cores on HECToR. This work has been submitted to the Conquest code repository and is used within groups at UCL and Imperial College.

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

The Block Overset Fast Flow high-order CFD Solver (BOFFS) is used to perform high resolution Large Eddy Simulations (LES) of turbulent flows for turbomachinery applications. The aim of this project was to enable more realistic turnaround times for grids with more than 10^5 points, by enabling more efficient use of HECToR. The project was allocated 4 months support at the June 2011 round and was supervised by Professor Paul Tucker of the Department of Engineering at the University of Cambridge. The work was performed by Lucian Anton, Ning Li and Phil Ridley of NAG, between February 2012 and July 2012.

By implementing asynchronous MPI calls for the inter block data transfers, the packing and unpacking of data buffers is now performed while each MPI process is waiting; this also enables grids with more complex block structures, and reduces the amount of time taken for the inter block communication by a third. Static arrays for all main variables were also replaced by allocatable versions to provide better memory utilisation. Furthermore, the OpenMP sections for the intra block computations in the tri-diagonal solver were improved by implementing a red-black decomposition. BOFFS can now use 8 threads per MPI task efficiently and up to 1.4 x speedup can be achieved compared with the original code. BOFFS is currently used on HECToR within the Computational Aeroacoustics Consortium.

6.7.18 VOX-FE Large Scale FE Bone Modelling on HECToR

VOX-FE is a finite element solver specifically designed for linear elasticity problems related 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 is relatively new on HECToR and the aim of this DCSE project was to enable the use of HECToR for an entire model of a femur or bone/implant interface. The project was allocated 4 months effort at the June 2011 round and was supervised by Professor Michael Fagan of the Department of Engineering at the University of Hull. The work was performed by Nick Johnson of EPCC, between February 2012 and September 2012.

By using the parallel netCDF-HDF5 libraries, file sizes were reduced by up to a factor of 190, which in turn reduced I/O time by up to a factor of 7. Use of netCDF also allows files to be made self-describing and portable between systems, along with offering good opportunities for compression. The improved I/O shows up to a 90 x speedup over a single core. VOX-FE is currently used for Medical and Biological applications at the University of Hull.

6.7.19 Micromagnetic modelling of naturally occurring magnetic mineral systems

This is the second DCSE project concerning the MicroMag application. The specific objectives of this work were to develop the existing parallel implementation by improving the initialisation step and the load balancing for the unstructured data. The project was allocated 6 months support at the June 2011 round and was supervised by Professor Wyn Williams of the School of Geosciences at the University of Edinburgh. Chris Maynard of EPCC performed the work from February to August 2012, and Paul Graham of EPCC took over the work from September to December 2012.

An improved method for the parallel construction of the matrix for the linear solver was developed for PETSc, in addition to more effective load balancing for the main solver from the use of better unstructured data partitioning. This work has enabled large-scale simulations with increased system size and precision, which facilitates more realistic theory to compare with experimental observation. MicroMag is currently used within Professor Williams' group at Edinburgh.

6.8 Ongoing DCSE Projects

6.8.1 NUMA-aware domain decomposition using Space-Filling Curves and other mesh decomposition techniques

The overall objective of this project is to improve the scaling behaviour of Fluidity. This will be achieved by implementing memory access techniques to exploit the mixed-mode OpenMP/MPI code within the finite element assembly routines. Different strategies for mesh entity renumbering and colouring for extracting low level concurrency will be implemented; first touch aware array initialisation algorithms will also be developed to ensure maximum benefit is obtained from NUMA architectures. This fourth Fluidity (ICOM) DCSE project was allocated 6 months support at the November 2011 round and was proposed by Dr Jon Hill of the Department of Earth Science and Engineering at Imperial College. Mark Filipiak of EPCC started the work in February 2012 and is scheduled to complete in January 2013.

6.8.2 Preparing DL_POLY_4 for the Exascale

Following on from four previous DCSE projects for improving DL_POLY_4, the first objective of this work is to implement a full mixed-mode OpenMP/MPI version of the code to exploit the shared memory features of HECToR. The second is to enable billion atom simulations by implementing a 64-bit integer representation within the code. This project was allocated 12 months effort at the June 2011 round and was proposed by Dr Ilian Todorov of STFC. Ian Bush of NAG started the mixed-mode work in November 2011. Asimina Maniopoulou, also of NAG, started work on the 64-bit integer representation in December 2011. The project is scheduled to complete early January 2013.

6.8.3 Optimising the Parallelisation of a Harmonic Balance Navier-Stokes CFD Code for the Ultra-rapid Analysis of Wind Turbine, Turbomachinery and Aircraft Wing Unsteady Flows

This project is to improve the parallel performance of the COSA CFD code. COSA is based on the compressible Navier-Stokes model for unsteady aerodynamics and aeroelasticity of fixed structures, rotary wings and turbomachinery blades. COSA may be used as a general-purpose time-domain solver or as a frequency-domain solver, utilising harmonic balance (HB) methods. The aim of this work is to improve the parallel efficiency of the HB solver for at least 512 cores, in both its pure MPI and MPI/OpenMP hybrid implementations. This project was allocated 6 months support at the November 2011 round and was proposed by Dr M. Sergio Campobasso of the School of Engineering at the University of Glasgow. Adrian Jackson of EPCC started the work in March 2012 and is scheduled to complete in February 2013.

6.8.4 Improved Global High Resolution Chemistry-Aerosol Modelling for Climate and Air Quality

TOMCAT-GLOMAP is a coupled chemical transport model and global model. This is the third TOMCAT-GLOMAP DCSE project and was allocated 6 months effort at the November 2011 round. The work was proposed by Professor Martyn Chipperfield of the School of Earth and Environment at the University of Leeds. The aim of this work is to enable global spatial resolutions of 1.0^{o} for more realism and improved ease of use. The existing Master I/O model will be replaced by MPI-IO, using a row-master-IO approach where each MPI task will read in a patch of data. Furthermore, general code optimisations will be performed to parts of the code previously identified. Mark Richardson of NAG began working on the project in April 2012 and is scheduled to finish in March 2013.

6.8.5 Software Framework to Support Overlap of Communications and Computations in Implicit CFD Applications

Following on from two previous DCSE projects for the Incompact3D complex fluid-flow solver and 2DECOMP&FFT library, this work will implement an overlap of communication and computation (OCC) algorithm within the underlying 2DECOMP&FFT library. This will be achieved by using non-blocking MPI collectives (e.g. MPI_ALLTOALL). The end result will enable more efficient performance on multi-core architectures for Incompact3D and other CFD applications that use 2DECOMP&FFT. This project was allocated 5 months support at the November 2011 round and was proposed by Professor John C. Vassilicos of the Department of Aeronautics at Imperial College. Ning Li of NAG began working on the project in April 2012 and is scheduled to finish in January 2013.

6.8.6 Microiterative QM/MM Optimisation for Materials Chemistry

This project was allocated 9 months effort at the November 2011 round and was proposed by Prof C. Richard Catlow of the Department of Chemistry at UCL. The aim of the work is to reduce the time to solution in Chemshell by implementing better structural determination algorithms within the DL_FIND component. Three established Quantum Mechanical/Molecular Mechanical (QM/MM) methods will be implemented: microiterative energy minimisation, microiterative transition state optimisation and microiterative reaction path optimisation. Tom Keal of STFC began working on the project in April 2012 and is scheduled to finish in March 2013.

6.8.7 Scalable coupling of Molecular Dynamics (MD) and Direct Numerical Simulation (DNS) of multi-scale flows Part 2

Following on from the first project for developing a coupled version of the Transflow (CFD) and Stream (MD) applications, this work will further improve the implementation for a fully coupled and scalable simulation of molecular dynamics (MD) with CFD. The first objective is to improve the load balance and scalability of the coupled applications, the second is to develop a user interface to facilitate the creation of other coupled applications. The project was allocated 6 months support at the November 2011 round. The work was proposed by Dr Tamer Zaki of the Department of Mechanical Engineering at Imperial College, who has also been co-ordinating the work since May 2012. The project is scheduled to complete at the end of February 2013.

6.8.8 Multigrid solver module for ONETEP, CASTEP and other codes

The ONETEP and CASTEP DFT codes are both widely used on HECToR and the overall objective of this project is to develop the functionality for a solvent model within both applications. This will be achieved by the implementation of a scalable multigrid Poisson solver module. The project was allocated 8 months effort at the November 2011 round and was proposed by Dr Chris-Kriton Skylaris of the Department of Chemistry at the University of Southampton. Lucian Anton of STFC began working on the project in April 2012 and is scheduled to complete the work in September 2013.

6.8.9 Seismic wave calculation at the Hz level for a general Earth

This project was allocated 5 months support at the April 2012 round. The SPECFEM3D_GLOBE package uses the (continuous Galerkin) spectral-element method to simulate three-dimensional global and regional seismic wave propagation. Current use of the code on HECToR is limited, due to numerical instabilities which occur when performing high-frequency simulations. These are caused by pervasive compiler optimisations, and the aim of this project is to resolve these and further improve performance by improving the initial mesh input operation. The work will then enable simulations on long-wavelength seismic phenomena, where the maximum frequency is around 0.01 Hz. The project was proposed by Dr James Wookey of the School of Earth Sciences at the University of Bristol. Paul Roberts of NAG began the work in July 2012 and is scheduled to complete in May 2013.

6.8.10 Performance enhancement and optimization of the TPLS and DIM two-phase flow solvers

The Two-Phase Level Set (TPLS) and Diffuse-Interface Method (DIM) codes are suites of two-phase solvers, which are used for three-dimensional Navier-Stokes problems. The objectives of this project are to:

- Optimise the MPI in TPLS by replacing the current point-to-points with non-blocking versions

- Replace the current successive over-relaxation pressure solver (for inverting Laplace's equation) with a hybrid OpenMP/MPI mixed-mode preconditioned conjugate-gradient routine

- Implement MPI in DIM for improved interface capturing

The work will enable TPLS and DIM to use HECToR to simulate flows on grids with at least 24 million points. The project was proposed by Dr Prashant Valluri of the School of Engineering at the University of Edinburgh and was allocated 8 months effort at the April 2012 round. David Scott of EPCC began working on the project in July 2012 and is scheduled to complete in April 2013.

6.8.11 Opening up HPC to the Discrete Element Method User Community

Granular LAMMPS is a scalable code that may be used for Discrete Element Modelling (DEM). One such application of DEM is to simulate granular response in materials for civil and geotechnical engineering applications. However current simulations are restricted due to lack of required functionality in the HPC variants of these applications. The aim of this project is to enhance the functionality of Granular LAMMPS by implementing new code for two sets of boundary conditions and a new contact model for bonding between two grains. The project was proposed by Dr Catherine O'Sullivan of the Department of Civil and Environmental Engineering at Imperial College and was allocated 7 months support at the April 2012 round. George Marketos, who is a member of Dr O'Sullivan's group, began the work in October 2012 and is scheduled to complete in April 2013.

6.8.12 Improving CONQUEST to allow ab initio molecular dynamics calculation on 100,000+ atoms

This is the fourth DCSE project for Conquest. The aim is to develop dynamic load balancing for better efficiency when moving atoms. The improved load balancing will be implemented for both the initial distribution of atoms and integration grid points to processes, and their dynamic redistribution during atom movement. These changes will enable significant gains in efficiency, permitting *ab initio* molecular dynamics calculations on systems with 10,000–100,000 atoms and beyond. The project was allocated 6 months effort at the April 2012 round and is supervised by Dr David Bowler of the Department of Physics and Astronomy at UCL. Lianheng Tong, also of UCL, began the work in October 2012 and is scheduled to complete March 2013.

6.8.13 Enhancement of high-order CFD solvers for many-core architecture

This is the second DCSE project for the Block Overset Fast Flow Solver (BOFFS) and was proposed by Professor Paul Tucker of the Department of Engineering at the University of Cambridge. The project was allocated 8 months support at the September 2012 round. The first aim of the project is to update the data decomposition in BOFFS to enable improved scalability; the second is to improve the hybrid (OpenMP/MPI) parallel performance of the related NEAT code. Phil Ridley of NAG started the work in December 2012 and is scheduled to complete in September 2013.

6.8.14 Massively Parallel Computing for Incompressible Smoothed Particle Hydrodynamics (ISPH)

ISPH is an efficient fluid solver which uses the incompressible smoothed particle hydrodynamics method for violent free-surface flows on offshore and coastal structures. The overall objective of this project is to optimise the ISPH code for at least 1024 cores to enable simulations with 10^8 particles. The optimisations will target:

- The mapping functions for particles data migration with one-sided MPI and the Zoltan library
- The neighbour searching algorithm by implementing a preconditioned dynamic vector method
- The pressure Poisson solver by improving the existing matrix representation which uses PETSc data structures

The project was proposed by Dr Benedict Rogers of the School of Mechanical, Aerospace and Civil Engineering at the University of Manchester. It was allocated 9 months effort at the September 2012 round. Xiaohu Guo of STFC started the work in December 2012 and is scheduled to complete in September 2013.

6.9 New DCSE Projects Starting 2013

6.9.1 Parallelisation of adaptive Kinetic Monte Carlo code

This project was proposed by Dr John Purton of STFC and was allocated 9 months support at the September 2012 round. David Gunn also of STFC will start the work in January 2013 with a scheduled completion in September 2013. The objective is to develop a hybrid OpenMP/MPI parallel version of the adaptive Kinetic Monte Carlo (aKMC) code. This will be achieved by distributing the independent saddle point searches over multiple nodes and developing a hybrid OpenMP/MPI parallelisation for the multiple force/energy evaluations required for each saddle point search. The serial aKMC code is currently used for simulations on radiation damage in materials and this DCSE project will enable systems with 200,000 atoms to be studied on HECToR.

6.9.2 Implementing OpenIFS on HECToR

OpenIFS is developed at the European Centre for Medium-Range Weather Forecasts (ECMWF); the code is already on HECToR but restricted to single process I/O, making it unsuitable for production runs. The aim of this project is to implement, configure and optimise the OpenIFS I/O framework software on HECToR. The project was proposed by Dr Grenville Lister of the Department of Meteorology at the University of Reading and was allocated 6 months effort at the September 2012 round. Mark Richardson of NAG will start the work in January 2013 with a scheduled completion in September 2013.

6.9.3 DL_POLY_4 algorithmic improvements

This is the fifth DL_POLY_4 DCSE project and was allocated 11 months support at the September 2012 round. The aims of the project are to potentially speed up the code by 10 to 30% for all users and reduce the time taken in calculations of certain expensive operations such as long-ranged Ewald and short ranged inter-molecular interactions. This will be achieved by implementing a Verlet neighbour list (VNL) calculation and symplectic time stepping routine. The work will be supervised by Dr Ilian Todorov of STFC, with Ian Bush of NAG performing the work on the VNL development along with a software development from STFC who will be implementing the symplectic time stepping routine. The work will being in January 2013 with a scheduled completion in September 2013.

6.9.4 Further Improving NEMO In Shallow Seas (FINISS)

The overall aim of this third NEMO DCSE project is to improve performance of the code on HECToR for shelf-sea studies. As the code was originally developed for vector architectures there is scope for improvement. The work will reduce redundant computations and communications beneath the seabed; facilitate more flexible partitioning strategies to improve the effective load balance; and implement loop-level avoidance of dry points. The project was proposed by Dr Stephen Pickles of STFC and was allocated 6 months effort at the September 2012 round. Andrew Porter, also of STFC, will start the work in January 2013 with a scheduled completion in August 2013.

6.9.5 Improving the scaling and performance of GROMACS on HECToR using single-sided communications

The GROMACS biochemical molecular simulation package has recently undergone a major restructuring for hybrid MPI/OpenMP. However, performance of the MPI point-to-points restricts scalability and the number of effective OpenMP threads allowed per MPI task. This project will replace the point-to-points with SHMEM equivalents to improve both scaling and parallel efficiency, enabling simulations for a million atoms which can efficiently utilise 16,384 cores. The project was proposed by Professor Jason Crain of the School of Physics and Astronomy at the University of Edinburgh and was allocated 6 months support at the September 2012 round. Ruymann Reyes Castro of EPCC will start the work in February 2013 with a scheduled completion of August 2013.

6.9.6 Adapting QSGW to large multi-core systems

This project was proposed by Professor Mark van Schilfgaarde of the Department of Physics at King's College and was allocated 6 months effort at the September 2012 round. Martin Lueders, Leon Petit and Lucian Anton of STFC will be performing the work. The work will begin in April 2013 with a scheduled completion of September 2013. Quasi-particle Self-consistent GW (QSGW) is an advanced method for *ab initio* electronic structure calculations, however it is comparatively more expensive than conventional methods. The current QSGW code has been developed for serial computation with up to 16 atoms per unit cell, this project will implement a parallel version to enable calculations on HECToR with 100 atoms per unit cell.

6.10 University Distributed CSE

EPSRC has funded a pilot scheme to offer DCSE to non-HECToR users (typically users of University facilities or the new regional HPC centres). The panel met in June 2012, and considered a total of fifteen applications requesting 153 months of effort. Of those, eight were recommended to receive a total of 82 months of effort. At the end of December all eight projects were running.

6.10.1 Parallelised Uncertainty Quantification using Differential Geometric Population MCMC

This project involves implementing an improved parallel Differential Geometric Markov Chain Monte Carlo (MCMC) sampler. The aim is to port an existing implementation which uses the MATLAB parallel toolbox to a scalable implementation in C with MPI. The existing MATLAB implementation is run on the Statistical Science Computing Cluster at UCL. Parts of this will be replaced iteratively with equivalent C routines using MATLAB's MEX interface, until there is a fully functional framework written entirely in C and MPI. The project was proposed by Professor Mark Girolami, of the Department of Statistical Science at UCL and was allocated 12 months support. Gary Macindoe, also of UCL, started the work in August 2012 and is scheduled to complete in July 2013.

6.10.2 Tight binding molecular dynamics on CPU/GPU clusters

The TBE code uses the tight binding approach, a simplified electronic structure method which is significantly faster than density functional calculations. TBE runs in parallel either in the k-point loop which applies largely in metals, or through the matrix diagonalisation using OpenMP or ScaLAPACK routines. The objectives of this project are to: implement a parallel Hamiltonian, charge and force calculation; improve the parallel diagonalisers; and develop a CUDA interface for MAGMA to replace ScaLAPACK. The project was proposed by Professor Anthony Paxton of the Department of Physics at QUB, and was allocated 12 months effort. Dimitar Pashov, also of QUB, started the work in July 2012 and is scheduled to complete in July 2013. The work will be carried out on the QUB Dell cluster.

6.10.3 Molecular Dynamics Simulation of Multi-Scale Flows on GPUs

The coupled OpenMM/OpenFOAM hybrid solver is used to simulate fast particle dynamics for multi-scale flow simulations. OpenMM has been developed for GPUs in order to solve the pairwise molecular force calculation. OpenFOAM employs an MPI-parallel decomposition for the flow solver. The aim of this project is to port more parts of the application to the GPU in order to speed up the calculation. The work was proposed by Professor Jason Reese, of the Department of Mechanical & Aerospace Engineering at the University of Strathclyde and was allocated 12 months support. Saif Mulla, also of the University of Strathclyde, started the work in July 2012 and is scheduled to complete in July 2013. The USE-HPC and ARCHIE-WeSt clusters at Strathclyde are being used for this work.

6.10.4 Scaling the Nektar++ Spectral/hp element framework to large clusters

Nektar++ is a tensor product based finite element package which has been designed to allow the numerical solution of partial differential equations using high-order discretisation. The code is implemented in C++ with MPI and the aim of this project is to improve scalability on the cx1 and cx2 clusters at Imperial College. This will be achieved by adding efficient parallel preconditioners and a new level of parallelism via the use of pthreads. The work was proposed by Professor Spencer Sherwin of the Department of Aeronautics at Imperial College and was allocated 12 months effort. Simon Clifford and David Moxey, also of Imperial College, started the work in September 2012 and are scheduled to complete in April 2013.

6.10.5 Expressive and scalable finite element simulation beyond 1000 cores

The finite element library DOLFIN is part of the FEniCS suite of applications for solving partial differential equations. It can be used to solve almost any partial differential equation, both linear and nonlinear. The aim of this work is to increase DOLFIN use on HPC resources by improving the scalability. This will be achieved by implementing hybrid OpenMP/MPI for the finite element assembly stage, parallel I/O, and parallel mesh refinement. The project was allocated 12 months support and was proposed by Dr Garth Wells of the Department of Engineering at the University of Cambridge. Chris Richardson, also of the University of Cambridge, started the work in August 2012 and is scheduled to complete in July 2013. The Darwin Cluster at Cambridge is being used for this work.

6.10.6 HPC Software for Massively Parallel Simulation of 3D Electromagnetic Nano-structures

OPTIMET (OPTIcal METamaterials) is based on the multiple-scattering (MSM) formalism for modelling electromagnetic wave interaction with assemblies of particles and material media. The aim of this project is to enhance the parallel functionality of OPTIMET for nano-device simulation and design. This will be achieved by extending the code capability to 3D geometries, implementing periodic boundary conditions and incorporating cubic optical nonlinearities. The work was proposed by Dr Nicolae Panoiu of the Department Electronic and Electrical Engineering at UCL and was allocated 12 months effort. Ahmed Al Jarro and Claudiu Biris, also of UCL, started the work in October 2012 and are scheduled to complete in August 2013. The Legion cluster at UCL is being used for this work.

6.10.7 Tensor Manipulation and Storage

Tensor Network Theory (TNT) provides efficient and highly accurate algorithms for the simulation of strongly correlated quantum systems. The TNT library is currently being developed by the group of Professor Dieter Jaksch of the Department of Physics at the University of Oxford. This project was allocated 10 months support and concerns the tensor storage part of the library. In particular, this will involve the optimisation and parallelisation of the core tensor functions. Chris Goodyer of NAG started the work in October 2012 and is scheduled to complete at the end of July 2013. The Theoretical Physics and Atomic and Laser Physics Departments' own cluster at Oxford is being used for the work. The overall aim is to enable the TNT library to make efficient use of HPC in order to be able to study two dimensional systems.

6.10.8 HPC implementation of Agent-Based Models for Cell Cultures

The capability to model complex networks of biochemical reactions over large populations of cells is vital for Synthetic Biology. Models for control and optimisation in this area are now employing agent-based techniques to capture the diversity in real cell cultures. This project was allocated 6 months effort and concerns the development of one such model, which is currently implemented in serial. The aim of the project is to develop an MPI parallel application for well-mixed, spatially located populations, as well as interacting cells. This will then enable larger and more complex bacterial populations to be studied. The work was proposed by Professor Declan Bates of the College of Engineering, Mathematics and Physical Sciences at the University of Exeter. Hugo Pinto, also of the University of Exeter, started the work in October 2012 and is scheduled to complete by April 2013. The Zen Astrophysics cluster at Exeter will be used for the development.

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	kAUs allocated	kAUs used	kAUs left
EPSRC Projects							
c01	Support of EPSRC/STFC SLA	EPSRC	Class1a	Dr Adrian Wander	54804	53099	1705
e01	UK Turbulence Consortium	EPSRC	Class1a	Dr Gary N Coleman	490370	167496	322874
e05	Materials Chemistry HPC Consortium	EPSRC	Class1a	Prof C Richard A Catlow	1139874	978490	161324
e10	GENIUS	EPSRC	Class1a	Prof Peter Coveney	257748	24212	233537
e68	Hydrogenation Reactions at Metal Surfaces	EPSRC	Class1a	Prof. Angelos Michaelides	50000	49888	113
e71	Simulating the control of calcite crystallisation	EPSRC	Class1a	Prof John Harding	130404	59373	71019
e76	HELIUM Developments	EPSRC	Class1a	Prof Ken Taylor	42522	40898	1624
e82	ONETEP: linear-scaling method on High Performance Computers	EPSRC	Class1b	Dr Peter Haynes	4853	4233	620
e85	Study of Interacting Turbulent Flames	EPSRC	Class1a	Dr N Swaminathan	8089	6286	1803
e89	Support for UK Car-Parrinello Consortium	EPSRC	Class1a	Dr Matt Probert	400100	371684	28416
e122	Multiscale Modelling of Magnetised Plasma Turbulence	EPSRC	Class1a	Dr Colin M Roach	150000	84886	65114
e124	Compressible Axisymmetric Flows	EPSRC	Class1a	Prof Richard D Sandberg	22888	23224	-344
e125	Full configuration interaction quantum monte carlo	EPSRC	Class1a	Dr Ali Alavi	168325	39006	129209
e126	Clean Coal Combustion: Burning Issues of Syngas Burning	EPSRC	Class1a	Prof Xi Jiang	25584	17521	8064
e127	Alternative drag-reduction strategies	EPSRC	Class1a	Prof Michael Leschziner	7000	4154	2846

Code	Project Title	Funding Body	Class	Principal Investigator	kAUs allocated	kAUs used	kAUs left
e128	Rate-Controlled Constrained Equilibrium	EPSRC	Class1a	Dr Stelios Rigopoulos	7092	6692	400
e129	Novel Hybrid LES-RANS schemes [ICL]	EPSRC	Class1a	Prof Michael Leschziner	7500	2627	4873
e130	Novel hybrid LES-RANS schemes [MAN]	EPSRC	Class1a	Prof Dominique Laurence	10500	10489	11
e141	A numerical study of turbulent manoeuvering-body wakes	EPSRC	Class1a	Dr Gary N Coleman	16350	7458	8892
e145	UK-SHEC Consortium	EPSRC	Class1a	Dr T.J. Mays	1192	496	693
e149	Fractal-generated turbulence and mixing: flow physics and	EPSRC	Class1a	Prof Christos Vassilicos	68083	52807	15276
e155	Modelling Cholesterol Deposits	EPSRC	Class1a	Dr David Quigley	10000	1309	8691
e156	Metal Conquest: efficient simulation of metals on petaflop	EPSRC	Class2b	Dr David Bowler	1600	332	1268
e158	Novel Asynchronous Algorithms	EPSRC	Class1a	Prof Nicholas J Higham	2500	651	1849
e159	Multi-layered Abstractions for PDEs	EPSRC	Class1a	Prof Paul Kelly	3816	36	3780
e160	Sustainable Software Generation Tools	EPSRC	Class1a	Prof Paul Kelly	20208	15801	4407
e161	Properties and Dynamics of Atomic Bose-Einstein Condensates	EPSRC	Class1a	Dr A White	69896	0	69896
e165	Multi-scale simulation of intense laser plasma interactions	EPSRC	Class1a	Dr Tony Arber	4872	0	4872
e174	3D instabilities in two-layer flows	EPSRC	Class1b	Dr Prashant Valluri	11495	3977	7518
e175	Fine-Scale Turbulence	EPSRC	Class1a	Prof Richard D Sandberg	50000	6755	43090
e179	Non-conservative dynamics	EPSRC	Class1a	Dr Daniel Dundas	87000	3071	83929
e182	Advanced Modelling of Two-Phase Reacting Flow	EPSRC	Class1a	Dr Edward S Richardson	8150	0	8150

Code	Project Title	Funding Body	Class	Principal Investigator	kAUs allocated	kAUs used	kAUs left
e183	Analysis of Processes in Hydrocarbon Fuel Droplets	EPSRC	Class1a	Prof Sergei Sazhin	8640	7944	696
e184	UK-RAMP	EPSRC	Class1a	Prof Ken Taylor	130500	1401	129099
e185	Chemistry of ceramic materials	EPSRC	Class1a	Prof John Harding	340000	6231	333769
e186	Step Change in Combustion Simulation	EPSRC	Class1a	Prof Kai Luo	70000	62556	7444
e187	IAGP: Integrated Assessment of Geoengineering Proposals	EPSRC	Class1a	Prof Piers Fosters	6030	657	5373
e191	CFD Analysis of Flight Dynamics	EPSRC	Class1a	Prof Kenneth Badcock	40500	4413	36087
e192	Physical properties of carbon nanotubes	EPSRC	Class1b	Dr Michael R C Hunt	10963	7203	3760
e202	Quantum Monte Carlo simulations	EPSRC	Class1a	Prof Matthew Foulkes	38345	14648	23697
e203	BeatBox - Realistic Cardiac Simulations	EPSRC	Class1a	Prof Vadim Biktashev	4500	4202	298
e204	Rare Events via Parallel Forward Flux Sampling	EPSRC	Class1a	Dr Rosalind Allen	5000	1645	3355
e206	FLAME Agent-Based Simulation Framework	EPSRC	Class1a	Prof Christopher Greenough	410	1	410
e207	Developing DL_POLY Molecular Dynamics Simulation code	EPSRC	Class1a	Dr Kostya Trachenko	25858	18040	7817
e213	Condensation/Evaporation Heat Transfer in Micro/Nanochannels	EPSRC	Class2a	Dr Huasheng Wang	400	10	390
e220	Study of interacting turbulent flames 2	EPSRC	Class1a	Dr N Swaminathan	26122	17022	9100
e223	Numerical modelling of aorta dissection	EPSRC	Class2a	Prof. Xiaoyu Luo	300	0	300
e226	Novel Vibrational Spectroscopic Techniques	EPSRC	Class1a	Dr Andrew D Burnett	1032	0	1032
e227	OPL	EPSRC	Class2a	Dr Radhika R. S. Saksena	50	46	4

Code	Project Title	Funding Body	Class	Principal Investigator	kAUs allocated	kAUs used	kAUs left
e228	Development of the potential of doped metal-oxide nanotubes	EPSRC	Class1a	Dr Gilberto Teobaldi	20218	701	19517
e229	DTC in Complex Systems Simulations	EPSRC	Class1a	Prof Jonathan W Essex	50000	10331	39669
e235	Modelling offshore wind	EPSRC	Class1b	Prof Simon Watson	2100	539	1561
e237	Simulating Coupled Protein Folding and Nucleic Acid Binding	EPSRC	Class2a	Dr Christopher Baker	400	399	1
e240	MicroMag	EPSRC	Class2b	Prof Wyn Williams	800	314	486
e241	Potential Energy Surfaces for Bio-molecular Simulations	EPSRC	Class1a	Dr Lorna Smith	500	1	499
e242	Study of the Green Fluorescent Protein Fluorophore	EPSRC	Class2a	Dr Garth Jones	400	0	400
e243	Tailored Structures for Orthopaedic Implantations	EPSRC	Class2a	Dr Carmen Torres-Sanchez	400	0	400
e244	VOX-FE: Large Scale FE Bone Modelling on HECToR	EPSRC	Class2b	Prof Michael Fagan	800	8	792
e245	Parallelisation of a harmonic balance NS solver	EPSRC	Class2b	Dr Sergio Campobasso	800	813	-13
e246	Numerical simulation of capillary blood flow	EPSRC	Class2a	Dr Ellak Somfai	400	0	400
e247	Tool development for multiscale protein folding simulations	EPSRC	Class2a	Dr Robert Best	400	221	179
e248	Testing of a Distributed Coordinate Descent Method	EPSRC	Class2a	Dr Peter Richtarik	400	102	298
e249	Feedback flow control for reducing the aerodynamic drag	EPSRC	Class1b	Dr Aimee Morgans	9860	2961	6899
e254	Ceramic Composites for Fusion Power	EPSRC	Class1b	Prof Sergei Dudarev	8371	501	7870
e255	Turbulent Drag Reduction	EPSRC	Class2a	Dr Pierre Ricco	400	0	400
e256	Hybrid simulation on heat transfer	EPSRC	Class2a	Dr Huasheng Wang	300	38	262

Code	Project Title	Funding Body	Class	Principal Investigator	kAUs allocated	kAUs used	kAUs left
e257	Global stability and sensitivity of fuel injectors	EPSRC	Class1b	Dr. Matthew P Juniper	1728	1321	407
e258	Morphology and electronic props of semiconducting polymers	EPSRC	Class1b	Prof Alessandro Troisi	5650	4281	1369
e259	DNS of multi-species fuel combustion	EPSRC	Class1b	Dr N Swaminathan	31505	13644	17861
e260	Microscopic gas diffusion-reaction model	EPSRC	Class1a	Dr Jochen Blumberger	4940	1918	3022
e261	Expressive and scalable finite element simulation	EPSRC	Class2b	Dr Garth Wells	800	1	799
e262	MC simulations of semiconductor nanostructures	EPSRC	Class2a	Prof Ian Galbraith	300	2	298
e263	Modelling the Elastic and Moisture Barrier Properties of Skin	EPSRC	Class1b	Dr Rebecca Notman	40100	4051	36049
e264	Metabolic efficiency in neurons with extended morphology	EPSRC	Class2a	Dr Biswa Sengupta	300	118	182
e265	HPC for the Discrete Element Method User Community	EPSRC	Class2b	Dr Catherine O'Sullivan	800	0	800
e266	Thermal and Reactive Flow Simulation on High-End Computers	EPSRC	Class1a	Prof Kai Luo	226800	0	226800
e267	Simulating free-surface flow and fluid-structure interaction	EPSRC	Class2a	Dr Ido Akkerman	300	0	300
e268	Modelling of marine renewable energy farms	EPSRC	Class1a	Dr Bjoern Elsaesser	10000	0	10000
e269	Atomic data for fusion diagnostics	EPSRC	Class1b	Dr Catherine C A E Ramsbottom	4870	0	4870
e270	Turbulent mass transfer at high Schmidt number	EPSRC	Class1b	Dr Maarten van Reeuwijk	17000	0	17000
e271	Cloverleaf: preparing hydrodynamics codes for exascale	EPSRC	Class1b	Dr Stephen Jarvis	20430	0	20430
e272	TOUCAN: TOwards an Understanding of CAtalysis on Nanoalloys	EPSRC	Class1a	Prof Roy L. Johnston	187000	0	187000
j01	JST	EPSRC	Class1a	Dr Andrew R Turner	71991	22248	49650

Code	Project Title	Funding Body	Class	Principal Investigator	kAUs allocated	kAUs used	kAUs left
STFC Projects							
p01	Atomic Physics for APARC	STFC	Class1a	Dr Penny Scott	10003	7630	2373
NERC Projects							
n01	Global Ocean Modelling Consortium	NERC	Class1a	Dr Andrew C Coward	255691	192952	62739
n02	NCAS (National Centre for Atmospheric Science)	NERC	Class1a	Dr Grenville GMS Lister	716228	529374	186854
n03	Computational Mineral Physics Consortium	NERC	Class1a	Prof John P Brodholt	558861	464028	94833
n04	Shelf Seas Consortium	NERC	Class1a	Dr Roger Proctor	157006	115156	41849
n99	NERC Training	NERC	Class1a	Dr Grenville GMS Lister	2	0	2
BBSRC Projects							
b09	Circadian Clock	BBSRC	Class1a	Prof Andrew A Millar	2000	1394	606
b100	Widening the BBSRC HPC User Base	BBSRC	Class1a	Dr Michael Ball	10000	641	9359
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	5587	106	5482
b14	Understanding supercoiling-dependent DNA recognition	BBSRC	Class1a	Prof Anthony Maxwell	42600	19740	22860
b15	Simulating bird and dinosaur footprints	BBSRC	Class2a	Dr Peter L Falkingham	300	0	300
Director's Time							
d11	NAIS	DirectorsTime	Service	Prof Mark Ainsworth	10000	5530	4471

Code	Project Title	Funding Body	Class	Principal Investigator	kAUs allocated	kAUs used	kAUs left
d15	HPC-GAP	DirectorsTime	Service	Dr David Henty	102	15	87
d21	GADGET	DirectorsTime	Service	Dr Adrian Jenkins	1000	19	981
d23	TEXT FP7	DirectorsTime	Service	Dr Mark Bull	1500	36	1464
d24	SBSI	DirectorsTime	Service	Dr Stephen Gilmore	2000	958	1042
d25	Code Scaling	DirectorsTime	Service	Dr Ken Rice	51500	17668	33832
d26	Guest Training Accounts	DirectorsTime	Service	Miss Elizabeth Sim	650	613	37
d29	Nu-FuSe	DirectorsTime	Service	Mr Adrian Jackson	1500	578	922
d30	PARTRAC	DirectorsTime	Service	Dr Mark Sawyer	200	124	76
d32	APOS-EU	DirectorsTime	Service	Dr Michele Weiland	1500	952	548
d35	PhD	DirectorsTime	Service	Dr Mark Bull	10	0	10
d36	Genome	DirectorsTime	Service	Dr Alan Gray	3460	0	3460
d37	CRESTA	DirectorsTime	Service	Dr Lorna Smith	21000	8236	12764
d38	Windfarm Simulation	DirectorsTime	Service	Mr Adrian Jackson	471	410	61
d39	NCSA access	DirectorsTime	Service	Mr Mark A Straka	1000	986	14
d40	Computational Chemistry at St Andrews	DirectorsTime	Service	Dr Herbert Fruchtl	2000	217	1784
d41	NPL Project	DirectorsTime	Service	Dr Ulrich Zachariae	45000	38739	6261
d42	Oxford Nanopore Technologies	DirectorsTime	Service	Dr Jayne Wallace	1170	734	436

Code	Project Title	Funding Body	Class	Principal Investigator	kAUs allocated	kAUs used	kAUs left
d43	ECDF	DirectorsTime	Class2b	Mr Tony Weir	13000	1698	11302
d44	Crucible	DirectorsTime	Service	Mr Iain A Bethune	1000	0	1000
d45	MSc in HPC 2012-2013	DirectorsTime	Service	Dr David Henty	1000	46	954
d46	Silicate melts with CP2K	DirectorsTime	Service	Mr Iain A Bethune	500	0	500
d47	PGS Project	DirectorsTime	Service	Dr Kevin Stratford	100	0	100
d49	Leiden	DirectorsTime	Service	Dr. Simon Portegies Zwart	200	0	200
External Projects							
e168	TEXT	External	Service	Dr Mark Bull	1500	80	1421
x01	HPC-Europa	External	Service	Dr Judy Hardy	44762	43446	1317
PRACE Projects							
pr1u0702	HYDROGEN-ILs	PRACE	Class1a	Dr Chris A Johnson	770333	1699	768634
pr1u0704	HIFLY	PRACE	Class1a	Dr Chris A Johnson	8450	8499	-48
pr1u0705	TanGrin	PRACE	Class1a	Dr Chris A Johnson	14084	14042	42
pr1u0706	SIVE-2	PRACE	Class1a	Dr Chris A Johnson	14000	14411	-411
pr1u0804	FULLDRUG	PRACE	Class1a	Dr Chris A Johnson	15014	15063	-49
pr1u0805	NanoTherm	PRACE	Class1a	Dr Chris A Johnson	9009	9061	-51
pr1u0806	NELC	PRACE	Class1a	Dr Chris A Johnson	11408	555	10853

Code	Project Title	Funding Body	Class	Principal Investigator	kAUs allocated	kAUs used	kAUs left
pr1u0807	PARAMETER	PRACE	Class1a	Dr Chris A Johnson	6720	277	6443
pr1u0808	PIPETURB	PRACE	Class1a	Dr Chris A Johnson	12600	13511	-911
pr1u0809	VIPforVPH	PRACE	Class1a	Dr Chris A Johnson	5346	1162	4184
pr1u0810	DrugEffluxMechanism	PRACE	Class1a	Dr Chris A Johnson	4965	0	4965
pr1u0902	ESM4OED	PRACE	Class1a	Dr Chris A Johnson	10272	949	9323
pr1u0903	ICREIMUTANTS	PRACE	Class1a	Dr Chris A Johnson	3766	13	3754
pr1u0904	MoMoGal	PRACE	Class1a	Dr Chris A Johnson	23976	140	23836
pr1u0905	MPI-FETI	PRACE	Class1a	Dr Chris A Johnson	9096	63	9033
pr1u0906	FORSQUALL	PRACE	Class1a	Dr Chris A Johnson	3156	1	3155
pr1u0907	GPCR4D	PRACE	Class1a	Dr Chris A Johnson	2260	0	2260