HECToR Annual Report 2013

01 January – 31 December 2013

Issue: 1.0

Contents

1 Introduction

This report covers the period from 1 Jan 2013 at 08:00 to 1 Jan 2014 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/2013.php

2 Executive Summary

2013 was another positive year for the HECToR Service. The highlights of the service over the year included:

- Average utilisation on the XE6 in 2013 was 101% of optimum, compared to 85% in 2012.
- There were only 8 technology-attributed service failures in 2013, as opposed to 12 in 2012. The overall MTBF was 1098 hours compared to 732 hours in 2012. Despite the fact that the hardware was reaching the end of service, the XE6 was very reliable.
- Downtime due to scheduled maintenance was kept to a minimum in 2013. The concept of using ‘at-risk’ periods for non-intrusive maintenance has worked well.
- A total of 6224 queries were handled in 2013 and the associated Helpdesk statistics for the year were excellent.
- The Research Data Facility (RDF) continued to play a key role in the service. Three “Data Mover Nodes” were configured on the RDF in 1Q13. These enable users to transfer data on/off the RDF to their home institution without the need to login to HECToR. In 2013, all user data which was resident on the HECToR Archive was copied across to the RDF.
- The ARCHER service was brought online in November. The transition from HECToR to ARCHER was designed in such a way to minimize user disruption. Using a single instance of the SAFE to manage both HECToR and ARCHER users is central to this. The transition of users has begun and will continue into March 2014 when the HECToR service will end.
- This last year has been a very successful one for the CSE service. The helpdesk continued to receive excellent feedback from users, as dealt with nearly 400 queries.

Training continued to be popular – there were over 2000 attendees at courses since the start of HECToR.

- The highlight of the CSE service continues to be the DCSE programme which is now coming to an end. In 2013, four projects were awarded IDC HPC Innovation Excellence Awards at the ISC and SC conferences in Leipzig and Denver.

- This year represented the final full year of service for HECToR and Cray Centre of Excellence (CoE) and Cray Exascale Research Initiative, Europe (ERI) teams took this opportunity to encourage and enable users to expand their applications on the next generations of supercomputers. This included collaborations with GPU Accelerator based systems like Titan, a Cray XK7, and progress for ARCHER, the Cray XC30 supercomputer that forms the next generation UK national service.

- The CoE and ERI have continued to build and maintain the close relationships between Cray's Application and Development teams and equivalent Cray application support teams in Europe and the wider world. This approach has been especially valuable for previewing technology to HECToR users that has been deployed at other sites and sharing experience and knowledge between sites to improve best practice.

3 Quantitative Metrics

3.1 Reliability

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

	1Q13	*2Q13*	*3Q13*	*4Q13*	*2013*
Incidents	28	40	36	39	143
Failures	1	3	3	1	8

The incidents above are primarily related to single node failures. Details on both the service failures and single node fails in 2013 can be found in Section 4.

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*	*1Q13*	*2Q13*	*3Q13*	*4Q13*	*2013*
Technology	Failures	1	3	3	1	8
	MTBF	2196	732	732	2196	1098
Service Provision	Failures	0	0	0	0	0
	MTBF	∞	∞	∞	∞	∞
External/Other	Failures	0	0	0	0	0
	MTBF	∞	∞	∞	∞	∞
Overall	**Failures**	**1**	**3**	**3**	**1**	**8**
	MTBF	**2196**	**732**	**732**	**2196**	**1098**

3.2 HECToR Utilisation

3.2.1 Phase 3 Utilisation

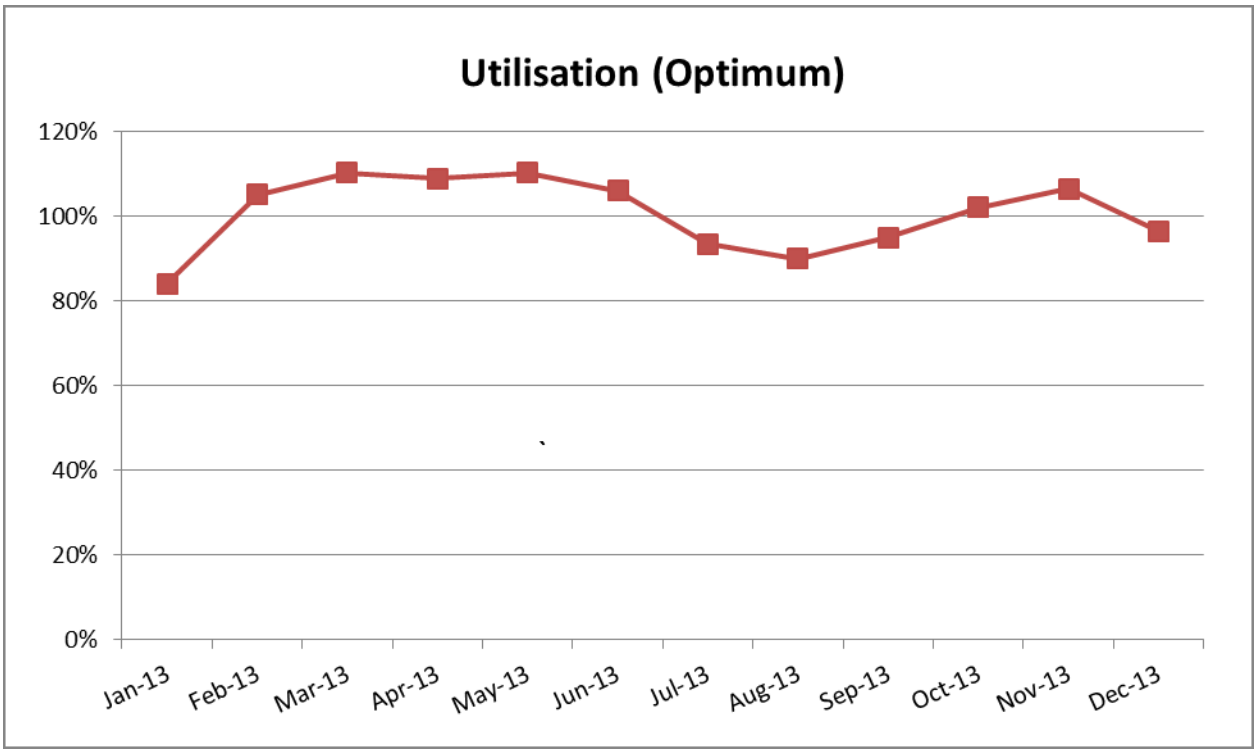

Overall utilisation of the XE6 in 2013 was 81% (101% of optimum), compared to 68% (85% optimum) in 2012.

3.2.2 HECToR Utilisation by Project Class

There are six main project classes on HECToR:

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

In 2013 the concept of Consortia allocations was implemented. 5 EPSRC and 4 NERC consortia changed from having very long AU allocations, to being allocated fixed %ages over a set time period. These allocations will be reviewed in March 2015 for EPSRC, and March 2014 for NERC.

The Consortia and Class1a projects transfer fully to ARCHER at the end of the HECToR Service. In the service overlap period they have access to both machines.

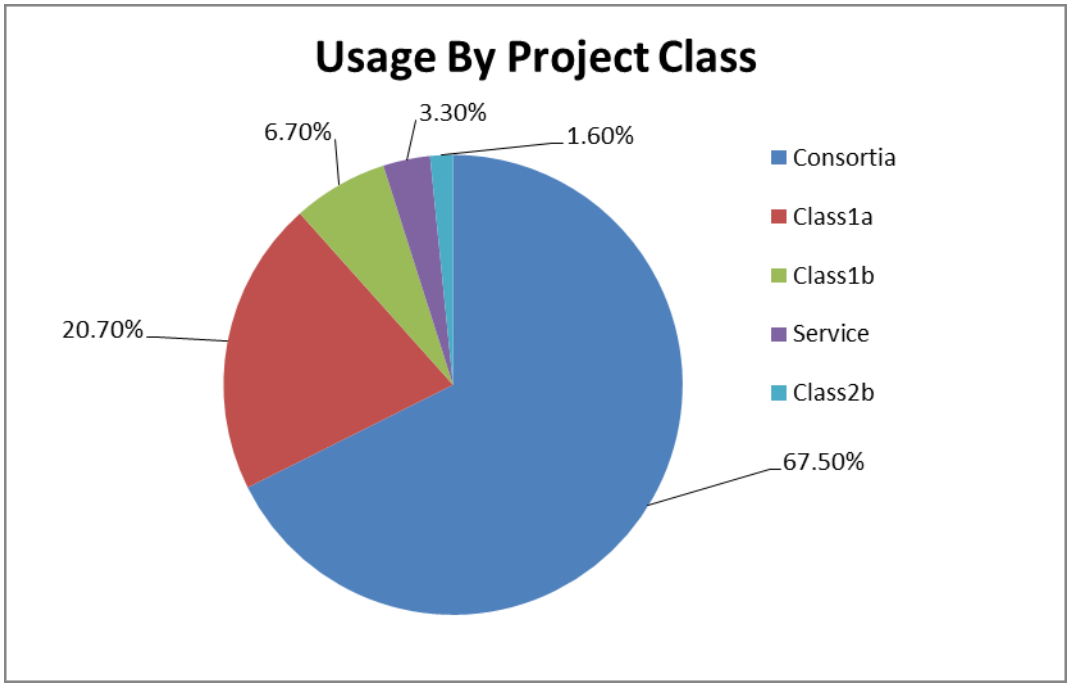

3.2.3 HECToR Utilisation by Project Code

As below, the main utilisation in 2013 came from the major EPSRC and NERC consortia as one would expect.

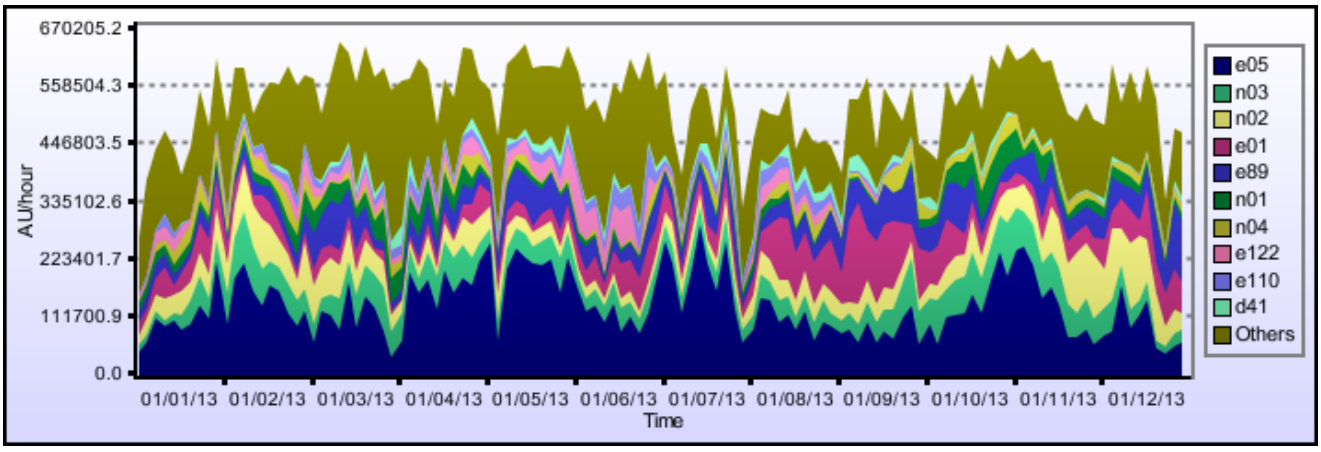

A total of 5,783,702.4 kAUs were available during this period. The utilization was as follows:

Project	Raw kAUs	Number of Jobs	Utilisation
y01	0.2	4	0.00%
y02	4.5	476	0.00%
y03	1,814.70	6,961	0.03%
y04	74.7	51	0.00%
y05	0	2	0.00%
y07	16.7	466	0.00%
z01	3,156.10	9,709	0.05%
z02	0.2	5	0.00%
z03	25,973.60	4,586	0.45%
z06	2.6	30	0.00%
z12	18.9	22	0.00%
Internal Total	**31,062.20**	**22,312**	**0.54%**
c01	20,502.20	5,928	0.35%
e01	402,032.50	19,617	6.95%
e05	1,191,159.10	123,201	20.60%
e10	45,931.50	7,141	0.79%
e42	0	4	0.00%
e63	0	1	0.00%
e68	22.9	10	0.00%
e71	30,894.80	7,879	0.53%
e72	0	4	0.00%
e74	0	5	0.00%
e76	2,798.10	207	0.05%
e82	1,155.70	315	0.02%
e89	368,525.80	76,945	6.37%
e104	6,071.80	1,174	0.10%
e107	1,341.10	204	0.02%
e108	22,928.70	1,300	0.40%

e109	0	1	0.00%
e110	105,558.00	6,390	1.83%
e122	119,896.60	8,471	2.07%
e124	0	7	0.00%
e125	23,413.60	1,263	0.40%
e126	12,446.80	107	0.22%
e127	453	8	0.01%
e130	45.7	27	0.00%
e135	0	2	0.00%
e136	397.7	96	0.01%
e137	0	1	0.00%
e138	0	1	0.00%
e141	2,568.20	730	0.04%
e142	0	20	0.00%
e145	683.7	538	0.01%
e148	0	1	0.00%
e149	18,298.60	282	0.32%
e155	7,062.20	393	0.12%
e156	121.6	153	0.00%
e158	1,041.50	154	0.02%
e159	353	1,178	0.01%
e160	9,923.70	1,135	0.17%
e165	0.1	9	0.00%
e174	3,736.40	2,826	0.06%
e175	34,398.20	2,361	0.59%
e179	12,173.60	538	0.21%
e182	1,058.40	181	0.02%
e183	3,083.30	34	0.05%
e184	41,756.10	1,982	0.72%
e185	4,857.90	163	0.08%
e186	14,505.10	1,194	0.25%
e187	736.6	845	0.01%
e191	2,663.90	131	0.05%
e192	8,295.70	356	0.14%
e202	48,541.60	1,218	0.84%
e203	263.4	167	0.00%
e204	5,466.90	809	0.09%
e207	29,453.30	70	0.51%
e213	21.2	30	0.00%
e215	0	2	0.00%
e220	0.1	7	0.00%
e228	6,051.80	569	0.10%
e229	25,887.60	4,950	0.45%
e231	16,423.30	845	0.28%
e235	1,446.10	980	0.03%
e240	54	149	0.00%
e245	74.6	34	0.00%
e248	83.8	167	0.00%
e249	39.3	17	0.00%
e251	0	1	0.00%
e254	20,802.40	844	0.36%
e256	1.4	22	0.00%
e257	1,037.10	56	0.02%
e258	3,518.30	24	0.06%
e259	73,759.60	133	1.28%

e260	5,003.70	1,633	0.09%
e261	426.9	210	0.01%
e262	19.3	82	0.00%
e263	65,547.40	2,226	1.13%
e264	164.2	3	0.00%
e265	507	88	0.01%
e266	52,520.40	3,296	0.91%
e268	0	2	0.00%
e269	9,207.30	301	0.16%
e270	41,998.00	169	0.73%
e271	45,899.20	385	0.79%
e272	38,718.70	1,384	0.67%
e273	523.5	299	0.01%
e274	8,262.10	2,102	0.14%
e275	257.3	204	0.00%
e276	516.8	186	0.01%
e277	609.8	438	0.01%
e278	476.8	52	0.01%
e279	410.2	146	0.01%
e280	35,767.90	3,302	0.62%
e281	19,068.00	3,449	0.33%
e282	14.7	104	0.00%
e283	427	30	0.01%
e284	3,592.90	394	0.06%
e285	1,559.50	53	0.03%
e286	5,426.20	107	0.09%
e287	1,617.90	180	0.03%
e288	0	8	0.00%
e289	0	1	0.00%
e290	215.8	401	0.00%
e291	630.1	306	0.01%
e292	4,997.00	11	0.09%
e294	6	13	0.00%
e296	3,265.50	978	0.06%
e297	670.7	2,363	0.01%
e298	389.7	156	0.01%
e301	60.6	12	0.00%
j01	2,019.50	172	0.03%
u03	0	1	0.00%
u09	0	1	0.00%
EPSRC Total	**3,106,616.80**	**311,855**	**53.71%**
n01	173,151.40	73,039	2.99%
n02	404,124.00	124,937	6.99%
n03	405,564.10	58,955	7.01%
n04	141,621.80	26,446	2.45%
u05	0	1	0.00%
u07	0	2	0.00%
NERC Total	**1,124,461.40**	**283,380**	**19.44%**
b09	1,143.70	239	0.02%
b12	0	2	0.00%
b14	22,493.80	1,585	0.39%
b15	514.2	162	0.01%
b16	4,836.40	66	0.08%
u04	0	1	0.00%
BBSRC Total	**28,988.10**	**2,055**	**0.50%**

p01	10,593.90	375	0.18%
STFC Total	**10,593.90**	**375**	**0.18%**
x01	387.5	310	0.01%
External Total	**387.5**	**310**	**0.01%**
b10	69.2	645	0.00%
d03	0	1	0.00%
d06	0	1	0.00%
d11	2,303.90	689	0.04%
d15	719.9	722	0.01%
d16	0	1	0.00%
d21	0	2	0.00%
d25	4,909.60	632	0.08%
d26	929.9	2,996	0.02%
d27	0	1	0.00%
d29	2,353.30	568	0.04%
d32	81.4	58	0.00%
d34	0	2	0.00%
d37	14,500.00	2,080	0.25%
d40	2,503.10	307	0.04%
d41	91,351.50	5,355	1.58%
d42	667.9	30	0.01%
d43	72,086.60	1,983	1.25%
d45	2,745.00	6,406	0.05%
d46	193	5	0.00%
d49	2.3	18	0.00%
d51	685	198	0.01%
d52	40.1	121	0.00%
d53	0	17	0.00%
d54	25.8	580	0.00%
d56	20,018.80	297	0.35%
e303	56.4	10	0.00%
i10	267.2	38	0.00%
Directors Time Total	**216,509.90**	**23,763**	**3.74%**
pr1u0002	895.9	319	0.02%
pr1u0702	12,442.00	325	0.22%
pr1u0804	0	5	0.00%
pr1u0806	10,850.70	190	0.19%
pr1u0807	24,435.70	1,063	0.42%
pr1u0901	329.8	33	0.01%
pr1u0902	16,899.20	410	0.29%
pr1u0903	5,066.70	768	0.09%
pr1u0904	10,924.60	270	0.19%
pr1u0905	1,331.50	914	0.02%
pr1u0906	0.1	10	0.00%
pr1u0907	5,824.40	62	0.10%
pr1u1004	8,172.20	157	0.14%
pr1u1008	10,667.60	131	0.18%
pr1u1009	848.8	171	0.01%
pr1u1010	9,160.20	44	0.16%
pr1u1011	7,527.20	4,949	0.13%
pr1u1012	9,983.50	217	0.17%
pr1u1108	0	3	0.00%
PRACE Total	**135,359.90**	**10,041**	**2.34%**
i01	227.5	120	0.00%

i02	160.4	373	0.00%
i03	24.3	93	0.00%
i04	673.4	229	0.01%
i05	303.8	19	0.01%
i06	66.2	29	0.00%
i07	480.9	89	0.01%
i09	1.4	9	0.00%
Industrial Total	**1,938.00**	**961**	**0.03%**
Total	**4,655,917.70**	**655,052**	**80.50%**

3.2.4 HECToR Utilisation by Application Area

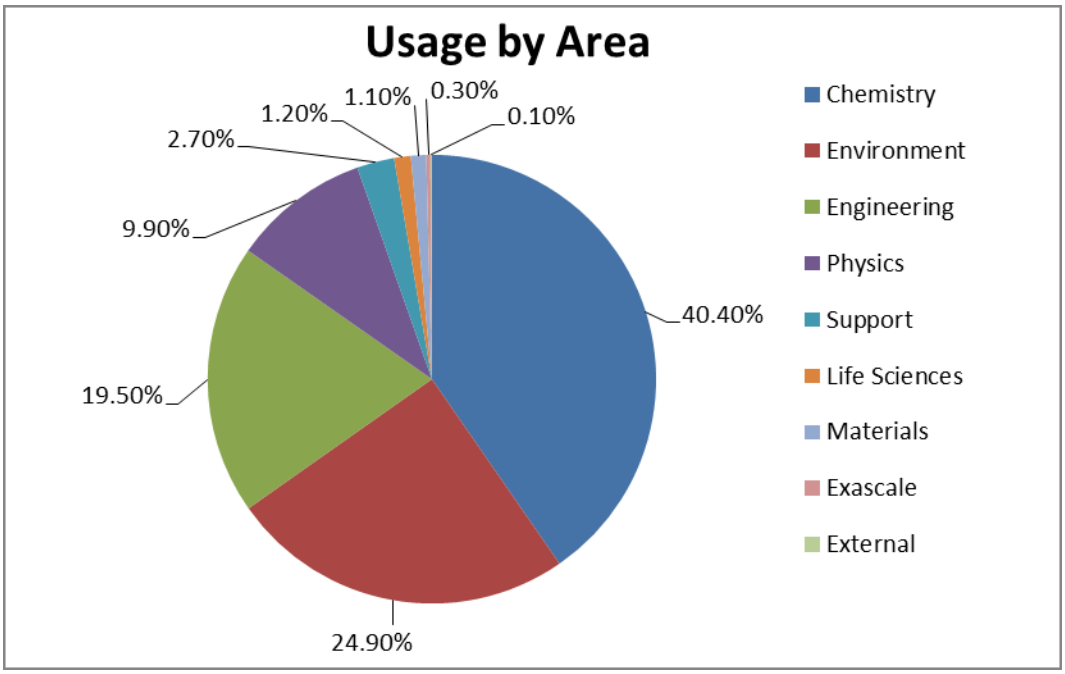

3.3 Helpdesk

A total of 6224 queries with a specified service metric were completed in 2013.

3.3.1 Helpdesk Targets

Metric	Pass	Total	Fraction	Target
All queries finished in 1 day	4185	4256	98.3%	97.0%
Admin queries finished in 1 day	4016	4079	98.5%	97.0%
Queries assigned in 30 min	4809	4826	99.6%	97.0%
Technical assessments in 10 days	67	69	97.1%	97.0%

This table does not include queries for the RDF, LMS or GPU. These have no performance metrics associated.

3.3.2 Queries by Service Metric

Service Metric	Queries	Percentage
Automatic	2876	46.2%
Automatic (RDF setup: non-service)	1264	20.3%
Admin	1203	19.3%
In-depth	501	8.0%
Technical	177	2.8%
No Metric (GPU,RDF)	134	2.2%
Technical Assessment	69	1.1%

3.3.3 Queries by Category

Query Category	Queries	Percentage
Create Directory	1163	18.7%
New User	661	10.6%
Set Group Quotas	591	9.5%
New Password	532	8.5%
Set User Quotas	400	6.4%
None	275	4.4%
Add to Group	257	4.1%
Disk, tapes, resources	234	3.8%
3rd Party Software	229	3.7%
Access to HECToR	225	3.6%
User Behaviour	175	2.8%
Compilers and System Software	163	2.6%
Batch System and Queues	160	2.6%
Other	137	2.2%
New Group	130	2.1%
Login, Passwords and ssh	109	1.8%
User Programs	108	1.7%
Make Reservation	85	1.4%

Join Project	78	1.3%
Remove Account	76	1.2%
Query Category	**Queries**	**Percentage**
SAFE	65	1.0%
Update Account	62	1.0%
Node Failure	53	0.9%
Courses	52	0.8%
Delete from Project	45	0.7%
Create Certificate	40	0.6%
Gpu	27	0.4%
Static Website	23	0.4%
Archive	17	0.3%
Performance and Scaling	14	0.2%
Delete from Group	12	0.2%
Network	9	0.1%
Grid	5	0.1%
Misc	12	0.1%

3.3.4 Queries by Handler Category

Handlers	Total	Automatic Non-Service	Automatic	Admin	In-depth	Technical	No Metric	Technical Assessment	%age
USL	1416			1043	150	128	95		22.8%
CSE	354			5	277		3	69	5.7%
OSG	4334	1264	2868	142	9	21	30		69.6%
Cray	120			21	65	28	6		1.9%

3.3.5 User Quality Tokens

22 positive and no negative quality tokens were received in 2013. A summary per project is below:

Project	***Positive Tokens***
e05	5
e10	4
e286	5
e277	4
d25	4
Total	**22**

3.4 Performance Metrics

Metric	TSL	FSL	January	February	March	April	May	June	July	August	September	October	November	December	Annual Average
Technology reliability (%)	85.0%	98.5%	99.6%	100.0%	100.0%	97.9%	100.0%	98.1%	88.6%	100.0%	99.9%	100.0%	100.0%	94.7%	98.2%
Technology MTBF (hours)	100	126.4	732.0	∞	∞	732.0	∞	366.0	366.0	∞	732.0	∞	∞	732.0	1098.0
Technology Throughput, hours/year	7000	8367	8705	8746	8690	8603	8724	8565	7724	8737	8780	8722	8716	8311	8483
Capability jobs completion rate	70%	90%	100.0%	100.0%	98.9%	100.0%	100.0%	100.0%	93.0%	100.0%	95.5%	100.0%	100.0%	100.0%	99.2%
Non in-depth queries resolved within 1 day (%)	85%	97%	98.1%	98.4%	99.0%	98.1%	98.5%	98.1%	97.8%	98.7%	98.4%	98.4%	98.2%	98.5%	98.3%
Number of SP FTEs	7.3	8.0	8.0	8.4	8.1	8.3	9.3	9.7	9.2	8.8	8.8	8.9	8.3	8.0	8.7
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	
Between TSL and FSL	
Below TSL	

4 Systems Hardware

4.1 HECToR Technology Changes

There were no major changes in the final year of service.

4.2 Severity-1 Incidents

4.2.1 Technology Failures

Cray supplied technology was responsible for 8 severity-1 incidents in 2013.

Incident	Date	Description
1	01-Feb-13	Compute blade power failure affecting system HSN
2	19-Apr-13	PBS batch system overload
3	11-Jun-13	esMS problems caused internal network lock
4	12 Jun-13	Cray XE6 cabinet blower control failure
5	06-Jul-13	/home filesystem cabinet circuit breaker trip
6	31-Jul-13	Lustre filesystem controller failure
7	05-Sep-13	Lustre filesystem OSS panic following disk timeouts
8	21-Dec-13	System Management Workstation filesystem failure

This is a breakdown of the failure categories:

- Two lustre filesystem related incidents
- One /home filesystem incident
- One cabinet blower control failure
- One compute blade fault which impacted the system High Speed Network
- One PBS batch system related failure
- Two System Management Workstation failures

The two lustre filesystem incidents were caused by hardware component failure in different storage devices. In both cases the faulty hardware was replaced and the system rebooted to restore the service to users.

A circuit breaker trip caused by an internal power supply fault occurred on one of the cabinets housing the /home filesystem. The faulty power supply was replaced and the service restored without the need for a full system reboot.

The speed control circuit for a Cray XE6 cabinet blower failed and required replacement before restoring the system to full service.

A compute blade power failure caused instability within the High Speed Network and required a system reboot to clear.

A PBS batch system reload was required when a user inadvertently submitted 175,000 jobs to the system causing a temporary overload of the PBS server. User limits have been enforced on the system to prevent similar problems from occurring in the future.

A System Management Workstation filesystem failure occurred which required a full system reboot to clear.

4.2.2 Other Failures

There were no failures attributed to either the Service Provision or External Sources in 2013.

5 Cray Centre of Excellence

5.1 Executive Summary

This year represented the final full year of service for HECToR and Cray Centre of Excellence (CoE) and Cray Exascale Research Initiative, Europe (ERI) teams took this opportunity to encourage and enable users to expand their applications on the next generations of supercomputers. This included collaborations with GPU Accelerator based systems like Titan, a Cray XK7, and progress for ARCHER, the Cray XC30 supercomputer that forms the next generation UK national service.

The CoE and ERI have continued to build and maintain the close relationships between Cray's Application and Development teams and equivalent Cray application support teams in Europe and the wider world. This approach has been especially valuable for previewing technology to HECToR users that has been deployed at other sites and sharing experience and knowledge between sites to improve best practice.

Members of both teams have also acted in advisory roles to the wider UK and European academic communities on the future of HPC hardware and software as perceived by Cray and in advising during the dCSE project review process. The ERI has also continued and extended its commitment to European HPC research with the continuation of the EU funded CRESTA (http://cresta-project.eu/) project and beginning the follow up EPiGRAM (http://www.epigram-project.eu/) exascale software research project. The following is a selection of user applications and projects have been involved with on the HECToR system and other Cray systems worldwide.

5.1.1 IFS/ECMWF

ECMWF uses the IFS model to provide medium range weather forecasts to its 34 European member states. Today's simulations use a global grid with a 16km resolution, but ECMWF expects to reduce this to a 2.5 km global weather forecast model by 2030 using an Exascale-sized system. To achieve this, IFS needs to run efficiently on a thousand times more cores. The CRESTA improvements have already enabled IFS to use over 200,000 CPU cores on Titan. Initial runs were done on the HECToR system up to 64,000 CPU cores. This is the largest number ever used for an operational weather forecasting code, and represents the first use of the pre-exascale 5km resolution model that will be needed in medium range forecasts in 2023. This breakthrough came from using new programming models to eliminate a performance bottleneck. For the first time, the Cray compiler was used to nest Fortran coarrays within OpenMP, absorbing communication time into existing calculations. Cray CoE and ERI staff assisted with porting and running the code to HECToR and Titan, and working closely with the users and Cray R&D to enable these developments. The graphs below show the improvement obtained through the use of coarrays and the raw scaling behaviour of the code at scale on both HECToR and Titan.

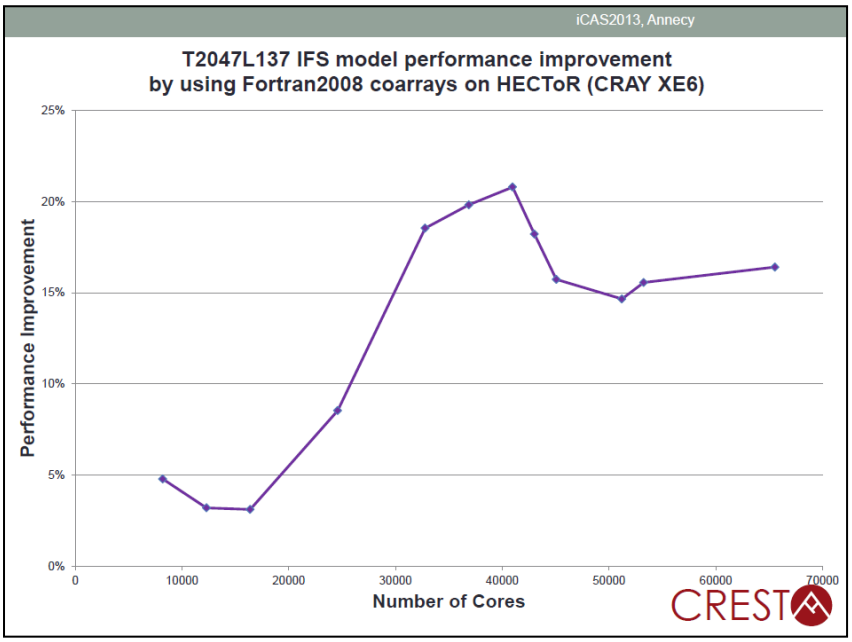

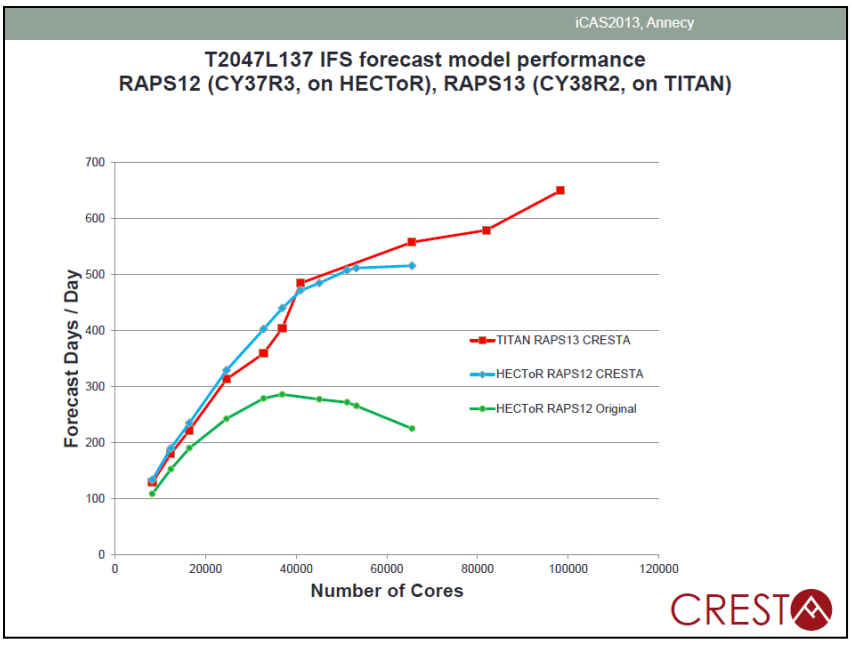

5.1.2 Computational Fluid Dynamics (CFD)/NEK5000

CFD applications are already some of the world's largest users of supercomputers and this importance is highlighted by CRESTA including two examples: OpenFOAM and NEK5000. CRESTA researchers have been using the OpenACC programming model to extend the existing NEK5000 code to portably and productively exploit accelerators. Adding only one OpenACC directive per thousand lines of Fortran code has already allowed a NEK5000 testcase to be efficiently scaled across more than 16,000 GPU nodes of Titan, with a near-threefold increase in performance compared to just using the CPUs.

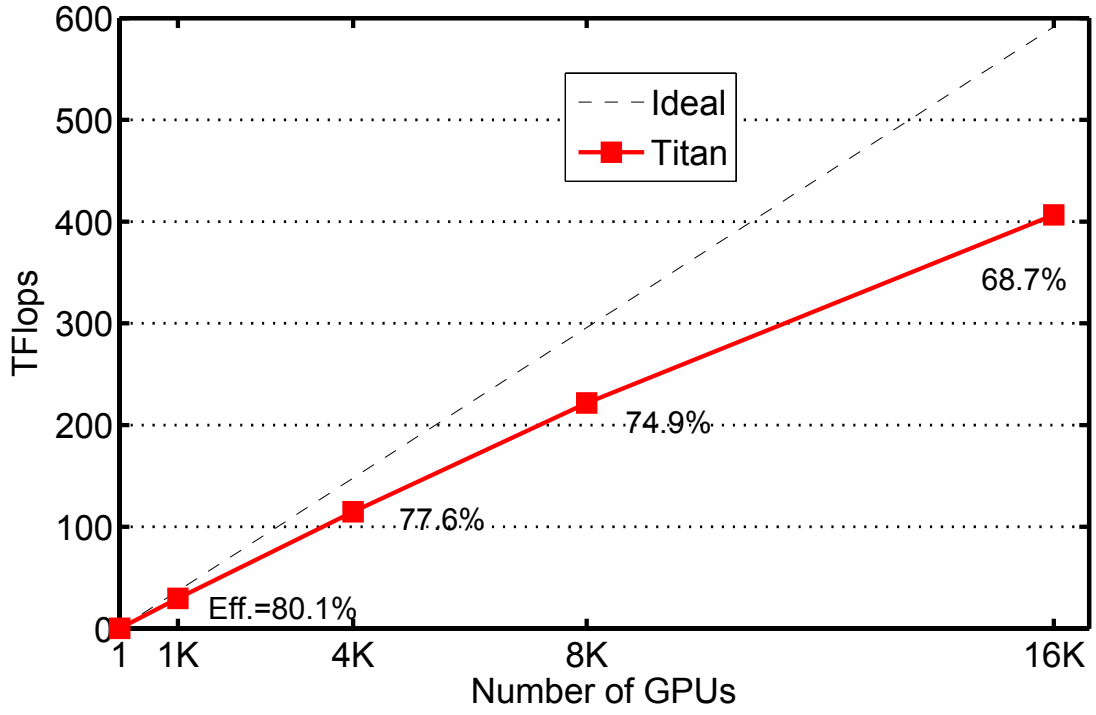

5.1.3 GROMACS

The GROMACS package simulates the molecular dynamics of millions of particles either in biochemical molecules (like proteins, lipids or nucleic acids) or non-biological systems, such as polymers. Work within the CRESTA project has focused on efficient, simultaneous exploitation of CPUs and GPUs with a scientific drive to understand the mechanism of membrane fusion in viruses. Cray CoE staff have been involved in using OpenACC to offload key kernels to a GPU. This is this then compared to performance obtained using a highly tuned CUDA implementation. This information is used to understand current bottlenecks, both in performance and programmability of both the directive based OpenACC approach and CUDA.

In parallel Cray staff at HECToR have been working closely with Cray compiler developers and GROMACS developers to support a Cray compiler build of GROMACS. GROMACS uses assembly intrinsics which are newly supported in the Cray compiler, so this has been a good test for this new functionality. This in turn benefits the GROMACS developers who will know have more choices in which compiler to use for their development work.

5.2 Seminars

Members of Cray's teams at HECToR have attended a variety of seminars and meetings through 2013, representing HECToR within the international HPC community.

- The CoE were invited to attend the “Sparse Linear Algebra Solvers for High Performance Computing Workshop” at Warwick University organized by AWE in July 2013.
- Members of the CoE participated in the “FEniCS ‘13” conference at Cambridge University where they presented on potential future directions for Cray HPC.
- Members of the CoE represented Cray at three dCSE Review Panels organized throughout the year
- The ERI were again invited to provide a Cray tutorial at Supercomputing 13 on OpenACC, an open standard for Accelerator based programming.
- Cray continued to support EPCC’s “Supercomputers and You” exhibition stand in its second year at the British Science Festival, the year in Newcastle-upon-Tyne.

5.3 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 around the world. These include:

- OpenACC/Cray XK7 Workshop: 18-19 Feb, CSCS, Switzerland
- Cray XC30 Workshop: 26 Feb – 1 March, CSC, Finland
- Cray XE6 Workshop: 16-19 April, Stuttgart, Germany
- OpenACC Workshop: 29-30 April, Stuttgart, Germany
- Cray XC30 Workshop: 14-17 May, CSC, Finland
- Cray XE6 PRACE PATC Workshop: 26-27 June, Edinburgh, UK
- Cray XC30 Workshop: 16-19 September, ECMWF, Reading, UK
- Cray XC30 Workshop: 16-19 September, ZIB, Berlin, Germany
- Cray XE6/XC30 Workshop: 28-31 October, Stuttgart, Germany
- PGAS2013: 3-4 October, Edinburgh, UK
- SC13 – OpenACC tutorial, SC13, Denver, USA

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 2013

This last year has been a very successful one for the CSE service. The helpdesk continued to receive excellent feedback from users, as we dealt with nearly 400 queries. Training continued to be popular and we have now had over 2000 attendees at our courses since the start of HECToR.

The highlight of the CSE service continues to be the DCSE programme which is now coming to an end. Originally designed just for HECToR users, in 2012 EPSRC provided additional funding to run a pilot for users of Regional and University services. The eight projects funded under this initiative ended this summer and all were extremely successful. Throughout the service, all completed DCSE projects have been assessed on how well they met their objectives and on the impact that they have had, by an independent panel established by EPSRC. The 73 projects assessed to date have scored, on average, 4.8 out of a maximum 5 points on the first criterion, and 4.2 out of 5 on the second. In addition, this year four projects were awarded IDC HPC Innovation Excellence Awards at the ISC and SC conferences in Leipzig and Denver. Chirag Dekate, Research Manager HPC/Technical Computing at IDC, said "NAG has done outstanding work in the HPC services arena. The HPC UF Steering Committee judged the HECToR DCSE Submissions to be compelling examples of advanced innovation with strong returns on investment, made possible by increasing the efficiencies of important applications kernels. IDC is delighted to present the NAG HECToR DCSE Team and their collaborators with their awards".

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.

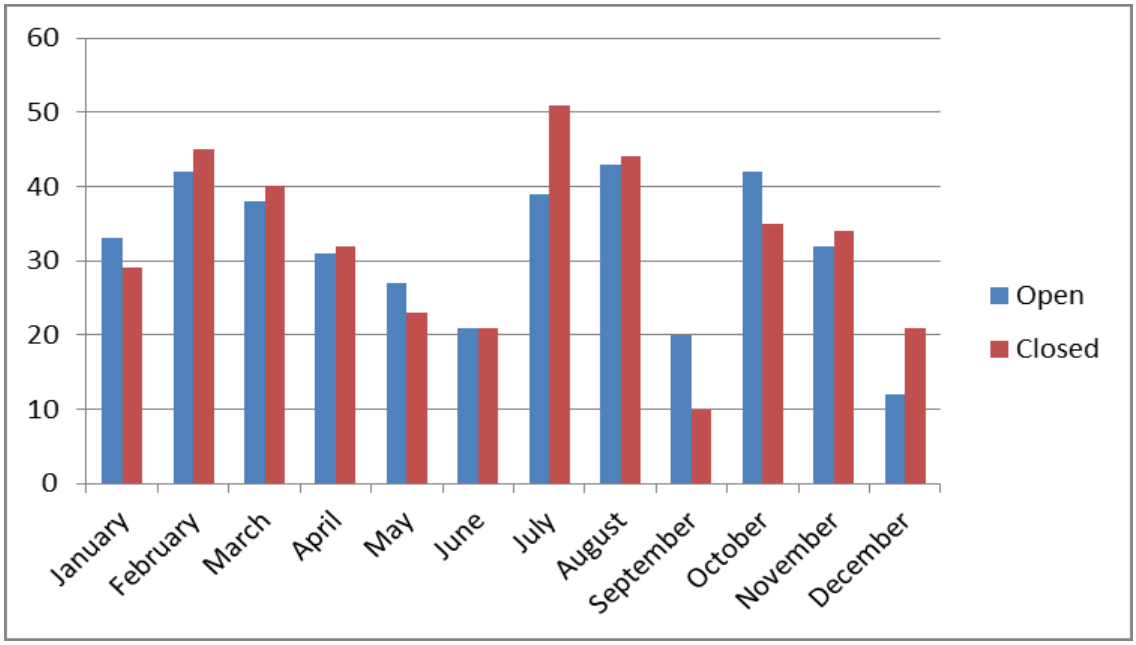

In cases where a technical query (as opposed to a request for a Technical Assessment) is resolved by the team (strictly speaking where the query is closed within the SAFE system by the CSE team), the user is invited to fill in a questionnaire giving feedback about his or her experience and satisfaction with the outcome. This year, 115 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?

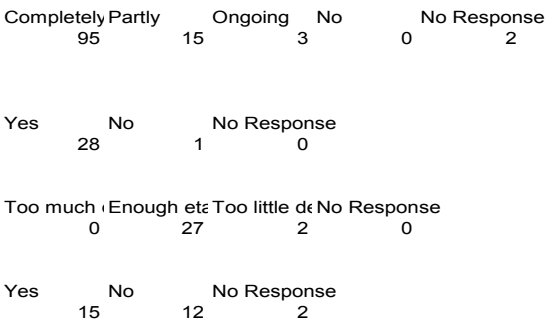

Q2: Were all communications replied to promptly and the text clear and understandable?

	Were all communications replied to promptly and the text clear and understandable?
Yes	112.5
No	2.5
No Response	0

To resize chart data range, drag lower right corner of range.

Q3: Were technical issues explained in sufficient detail?

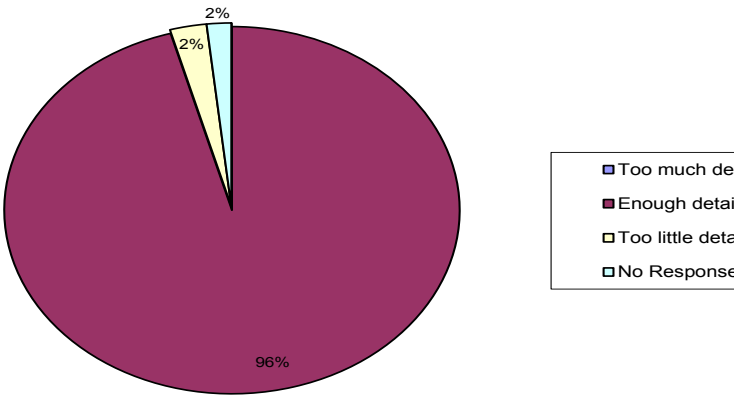

Q4: Did you use the HECToR CSE documentation to try and find a solution before submitting a query to the helpdesk?

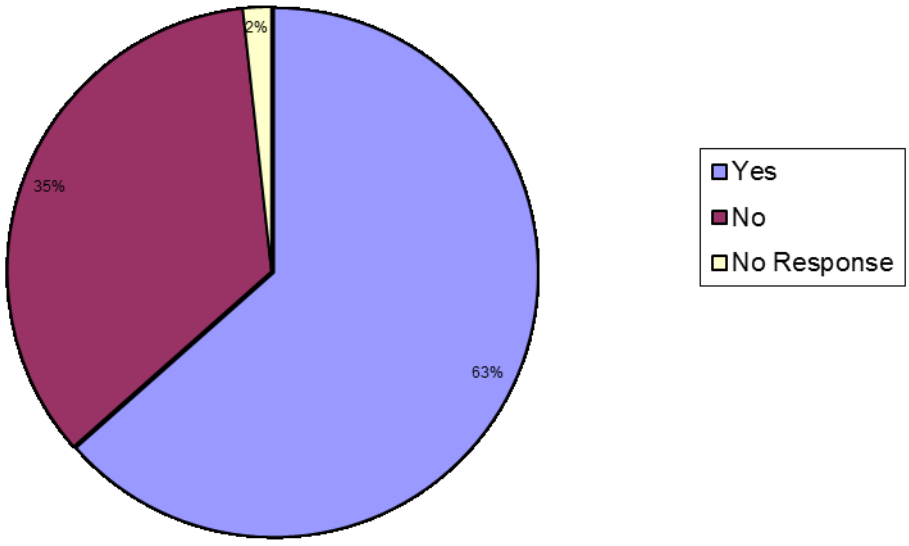

6.4 Training

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, 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 Science Centre, Daresbury Laboratory, Imperial College London, University College London, the Universities of Bath, Liverpool, Loughborough, Sheffield, Southampton, and Warwick, as well as at the NAG offices in Oxford and Manchester. A total of 73 days of training have been delivered to 468 attendees.

The complete list of courses offered in 2013 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.

- Accelerating Applications with CUDA
- An Introduction to CUDA Programming
- An Introduction to OpenCL Programming
- An Introduction to Unified Parallel C (UPC)
- Coarray Fortran
- Debugging, Profiling and Optimising
- Exploiting Parallel CASTEP on Large-scale HPC
- Fortran 95
- OpenMP
- Parallel IO
- Parallel Programming with MPI

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

There were no DCSE panel meetings during 2013, however twenty HECToR DCSE projects ended during the year. Since the start of the service the DCSE panel has allocated nearly 65 person-years of effort. In addition, nearly 7 further person-years were allocated to a pilot programme for University DCSE. This was to support HPC software development for local UK academic Tier 2 systems (with HECToR being at Tier 1). In total, eight projects were funded under this scheme, which ran from July 2012 until the end of July 2013.

6.7 DCSE Projects Completed in 2013

6.7.1 Preparing DL_POLY_4 for the Exascale

DL_POLY_4 is a general purpose package for classical molecular dynamics simulations. This fourth DCSE project for the code was allocated 12 months effort at the June 2011 call. The work was proposed by Dr Ilian Todorov of STFC. The first objective was to implement a full mixed-mode OpenMP/MPI version of the code, to make efficient use of more cores on HECToR (and future many-core architectures); and the second, to enable billion atom simulations by implementing a 64-bit integer representation within the code. Ian Bush of NAG performed the mixed-mode work between November 2011 and February 2013. Asimina Maniopoulou, also of NAG, performed the work for the 64-bit integer representation between December 2011 and March 2013.

For the mixed-mode work, the link cell and Verlet neighbour list routines were updated to include a second level of parallelism. OpenMP was also implemented for the evaluation of both the short and long ranged force terms, and the constraint force calculations. Furthermore, threaded parallelism was also implemented for the Ewald routines. A 64-bit integer kind representation was also introduced into the code. This was only for specific variables concerning the global index and, wherever possible, 32-bit kinds have been left to minimise the total memory footprint. The new code was successfully demonstrated for the cases of Argon, Sodium Chloride and Water. In each case at least 32,768 cores on HECToR were used, with up to 4,298,942,376 particles. The use of threads can now both increase the number of cores that DL_POLY_4 can exploit and also increase the overall performance. This can be up to a factor of 5-10, but is force field (and computational system) dependent. DL_POLY_4 is the fourth most heavily used single application on HECToR, consuming 5-10% of the total the machine resource.

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

Fluidity is a multi-scale, general purpose, finite-element CFD code. This fourth DCSE project for the code was allocated 6 months support at the November 2011 panel. The aim was to improve the ordering of topological and computational meshes and thus speed up the matrix assembly and matrix solver steps. Dr Jon Hill of the Department of Earth Science and Engineering at Imperial College supervised the work, which was performed by Mark Filipiak of EPCC between April 2012 and February 2013.

To improve the ordering of the nodes and elements of the mesh in each MPI process, a Hilbert space-filling curve (HSFC) method was implemented. This was achieved by using the HSFC from the Zoltan library, which was already used by Fluidity for mesh decomposition and load balancing across MPI processes. For a 3D backward-facing step simulation with 1,000,000 finite element nodes (grid points) the new mesh reordering now gives a 5% speed-up overall, although this is dependent on thread placement within the HECToR node. During 2013 Fluidity-ICOM used 19,149 kAUs on HECToR (0.6% of the overall machine and 0.3% of all jobs).

6.7.3 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

The COSA CFD code is based on the compressible Navier-Stokes model for unsteady aerodynamics and aeroelasticity of fixed structures, rotary wings and turbomachinery blades. This project was allocated 6 months effort at the November 2011 call. It was proposed by Dr

M. Sergio Campobasso of the School of Engineering at the University of Glasgow, and Adrian Jackson of EPCC performed the work between March 2012 and February 2013. The aim of the project was to improve the parallel efficiency of the COSA Harmonic Balance flow solver, by optimising both the MPI and the hybrid OpenMP and MPI implementations.

To improve the MPI point-to-point communications, COSA was updated to enable the use of halo data buffers, to reduce the number of messages. For the hybrid OpenMP and MPI version, the OpenMP within the code was improved for three main aspects: initialisation of the shared data to ensure cache coherency; single OpenMP parallel regions to reduce the work distribution overhead; and the placement of MPI within OpenMP regions to reduce any synchronisation overhead. The new code is now around 30% faster for 2048 processes and 40% faster with 512 processes. The overall runtime can now be reduced to half that of the original code, with a representative parallel efficiency of around 90%.

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

TOMCAT-GLOMAP is a coupled chemical transport model and global model. This was the third DCSE project to develop the code for HECToR, and it was allocated 6 months support at the November 2011 panel. The work was proposed by Professor Martyn Chipperfield of the School of Earth and Environment at the University of Leeds. Mark Richardson of NAG performed the work between April 2012 and March 2013. The aim of the work was to improve the performance of TOMCAT-GLOMAP to enable global simulations over several years at a higher resolution than 2.8 degree latitude x 2.8 degree longitude. This was achieved by refactoring the code to facilitate parallel I/O, and also by analysing high resolution simulations to provide recommendations for achieving the best scalability. This includes the number, combination, and placement of OpenMP threads and MPI tasks, in addition to recommended domain decompositions and compiler choices.

A detailed Performance analysis for a high resolution run showed that file access patterns for the supporting features of data collection prior to writing and data distribution post reading were inefficient. Two functions were identified and improved: reporting the altitude profiles and calculating the zonal means. The function for reporting the altitude profiles originally had an overhead of as much as 40 seconds. With the new version, this is now less than 0.05 second for a half second iteration. Similarly for the zonal means, the improvements have made this feature less expensive. This has enabled the researcher to increase the number of times when data is logged. Furthermore, a time series plot of such data may now reveal trends that might have been missed when using a lower frequency output. These developments are now in use on HECToR within the National Centre for Atmospheric Science (NCAS) group. The work was also presented at the Cray User Group Meeting in the US in May 2013.

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

The Incompact3D complex fluid-flow solver uses the highly scalable 2DECOMP&FFT library for parallelisation. This was the third DCSE project to develop these applications and it was allocated 5 months effort at the November 2011 call. The work was proposed by Professor John C. Vassilicos of the Department of Aeronautics at Imperial College. Ning Li of NAG performed the work between April 2012 and March 2013. The aim of the project was to reduce the time taken by the all-to-all communications in Incompact3D for large core counts by further developing the underlying 2DECOMP&FFT framework to enable the overlap of communications and computations (OCC).

The 2DECOMP&FFT library was developed to deal with flexible data layouts, enabling wider support for legacy applications and, in certain cases, better cache efficiency. To facilitate OCC, a set of low-level APIs were created for 2DECOMP&FFT and a high-level API was developed to enable multiple independent distributed FFTs to run from a single subroutine call. A sample application was also created to demonstrate the general use of fine-grain OCC for 3D distributed FFTs. The overall result is that, for representative sized problems, the OCC developments in the 2DECOMP&FFT library enable a performance gain of up to 15% to be achieved.

6.7.6 Microiterative QM/MM Optimisation for Materials Chemistry

ChemShell is a widely used computational chemistry environment. This was the second DCSE project to develop the application. The project was allocated 9 months support at the November 2011 panel and was proposed by Prof C. Richard Catlow of the Department of Chemistry at UCL. Tom Keal of STFC performed the work between April 2012 and March 2013. The aim of the work was to reduce the time to solution in Chemshell by implementing better structural determination algorithms in DL_FIND.

Three established Quantum Mechanical/Molecular Mechanical (QM/MM) methods were implemented: microiterative energy minimisation, microiterative transition state optimisation and microiterative reaction path optimisation. The microiterative (QM=macro cycles and MM=micro cycles) optimisation requires fewer expensive QM evaluations for convergence; on average between 5 and 12 times less than before. The developments will be made available to users in ChemShell version 3.6. During 2013, ChemShell used 1.8% of the overall allocation on HECToR, in 1.3% of all jobs.

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

The Transflow CFD code and the Stream-MD molecular dynamics code are two independent applications which may be used for coupled simulations. This work followed on from a previous DCSE project which had resulted in a successful implementation of the Continuum-to-Molecular Coupling and Embedding (CMCe) module, to interface the Transflow and Stream-MD applications for simulations over the molecular and continuum scales. The project was allocated 6 months effort at the November 2011 call and was proposed by Dr Tamer Zaki of the Department of Mechanical Engineering at Imperial College. Dr Zaki also coordinated the work between April 2012 and March 2013.

The load balancing and scalability was improved, together with the development of a new user interface. The weak scalability of TransFlow for a test case with a million grid points per MPI task can now achieve 90% parallel efficiency on HECToR and StreamMD with 3.3 million molecules can achieve 94% parallel efficiency. Therefore, the overall scalability with the CMCe module is now excellent for coupled DNS and MD simulations with millions of grid points and molecules, using several thousand of cores on HECToR. TransFlow, StreamMD and the CMCe module are used on HECToR within Dr Zaki's group under a resource allocation from the UK Turbulence Consortium.

6.7.8 Multigrid solver module for ONETEP, CASTEP and other codes

ONETEP is a linear-scaling code for quantum-mechanical calculations based on density-functional theory (DFT). ONETEP and CASTEP are both widely used on HECToR. This project was allocated 8 months support at the November 2011 panel and was proposed by Dr Chris-Kriton Skylaris of the Department of Chemistry at the University of Southampton. Lucian Anton of STFC performed the work between April 2012 and September 2013. The

aim of the work was to provide the functionality for a solvent model within both applications by developing a multigrid solver module.

A parallel hybrid (MPI and OpenMP) multigrid solver for the Poisson and Poisson-Boltzmann Equations was developed. 3D domain decomposition was used for the MPI parallelism, with the OpenMP threads operating on an adjustable grid block. The overall performance is almost 16 times faster than the previous solver in ONETEP. A representative problem for ONETEP with a global grid of 449x545x609 points demonstrated good scalability for up to 2048 HECToR cores (128 MPI processes, each with 16 threads). Furthermore, for a representative CASTEP problem with a global grid of 129x129x129 points, good scalability was demonstrated for up to 256 MPI processes. These developments will be made available to users in ONETEP version 3.6 and CASTEP version 7.2. During 2013 ONETEP used 31,915 kAUs on HECToR (1.8% of the overall machine and 2.0% of all jobs). CASTEP used 134,042 kAUs on HECToR (4.2% of the overall machine and 7.0% of all jobs).

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

The SPECFEM3D_GLOBE package uses the (continuous Galerkin) spectral-element method to simulate three-dimensional global and regional seismic wave propagation. This project was allocated 5 months effort at the April 2012 call and was proposed by Dr James Wookey of the School of Earth Sciences at the University of Bristol. Paul Roberts of NAG performed the work between July 2012 and June 2013. The aim of the project was to improve the performance of SPECFEM3D_GLOBE on HECToR, to enable simulations on long-wavelength seismic phenomena with a maximum frequency f~0.5 Hz.

SPECFEM3D_GLOBE was refactored to enable a higher level of compiler optimisation to be used. The testing and development of the refactored code was done using the Cray compiler (version 8.1.4). The initial mesh input operation was also improved by combining the mesh generation routine, xmeshfem3d with the main solver, xspecfem3D. Therefore, data is now passed between xmeshfem3d and xspecfem3D without any need for file access. This work has enabled SPECFEM_GLOBE to be used on HECToR for the CoMITAC five-year, ERC-funded, multidisciplinary project; to investigate the core-mantle interface (FP7/2007-2013).

6.7.10 Performance Enhancement and Optimization of the TPLS and DIM Two Phase Flow Solvers

Two-Phase Level Set (TPLS) and Diffuse-Interface Method (DIM) are suites of two-phase solvers, which are used for three dimensional Navier-Stokes problems. This project was allocated 8 months support at the April 2012 panel and was proposed by Dr Prashant Valluri of the School of Engineering at the University of Edinburgh. David Scott of EPCC performed the work between July 2012 and April 2013. The aim of the work was to enable TPLS and DIM to utilise HECToR efficiently, in order to simulate flows for grids with several million points.

Several redundant MPI barriers and some OpenMP collective operations were removed from TPLS, and some tuning of the OpenMP directives was performed. Also, non-blocking point-to-point communications were introduced to allow multiple communications to be overlapped and so reduce the latency overhead. For a test case with 11.2 million grid points, a 12% speedup was achieved. Furthermore, to improve performance of the most expensive area of the computation (the pressure calculation), calls to the PETSc 3.3 library were implemented to exploit conjugate-gradient and other Krylov based methods. Excellent strong and weak scaling of the code can now be achieved for 22.4 million grid points and 2048 cores, which is

up to a 54% improvement on the original code. Over 24,000 kAUs have been allocated for use with TPLS and DIM between May 2013 and April 2016.

6.7.11 Opening up HPC to the Discrete Element Method User Community

The Granular LAMMPS package is a scalable application 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. This project was allocated 7 months effort at the April 2012 call and was proposed by Dr Catherine O'Sullivan of the Department of Civil and Environmental Engineering at Imperial College. George Marketos who is a member of Dr O'Sullivan's group performed the work between October 2012 and April 2013. The aim of this project was to develop granular LAMMPS for added functionality by implementing new code for two sets of boundary conditions and a new contact model for bonding between two grains.

Granular LAMMPS was developed for moving and rigid stress controlled boundaries. For the addition of a membrane boundary condition, an algorithm based on the weighted Voronoi graph method was implemented by using the external package, Voro++. To develop a new contact model for the bonding between two grains, a non-linear elastic contact spring model (Hertz-Mindlin) was implemented and a new bonding model was added. The new work was demonstrated by performing the first known stress-controlled DEM simulation for geomechanics where more than 1,000,000 particles were compressed in a controlled manner. This required the use of 256 cores on HECToR for a run time of over 48 hours. Around 20,000 kAUs have been allocated for use with Granular LAMMPS before March 2014.

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

Conquest is a code which is based on *ab initio* Density Functional Theory (DFT). This fourth DCSE project for Conquest was allocated 6 months support at the April 2012 panel. The aim was to develop dynamic load balancing for better efficiency when moving atoms, to enable calculations on systems with 10,000-100,000 atoms and beyond. Dr David Bowler of the Department of Physics and Astronomy at UCL supervised the work which was performed by Lianheng Tong, also of UCL, between October 2012 and March 2013.

To improve the automatic partitioning and load balancing in Conquest, a non-cubic compact 3D Hilbert space space-filling-curve partitioning algorithm was implemented along with dynamical re-assignment. The partitioning process is now easier to perform and control. Furthermore, the computation time of certain simulations may now be reduced by 40%-90%, depending on the type of system. Additionally, extended Lagrangian Born-Oppenheimer molecular dynamics was also implemented. This allows more stable simulations to be performed with loose energy tolerances, which are also computationally cheaper. During 2012 there were 5,587 kAUs used for Conquest simulations on HECToR.

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

The GROMACS package is a scalable application which is used to perform molecular dynamics, mainly for biochemical simulation. This project was allocated 6 months effort at the September 2012 call and was proposed by Professor Jason Crain of the School of Physics and Astronomy at the University of Edinburgh. Ruymann Reyes Castro of EPCC performed the work between February 2013 and August 2013. The aim was to improve the performance of the MPI point-to-points to enable more efficient use of a higher number of cores per MPI task and further increase scalability.

The calls to the standard MPI two-sided point-to-point communication routines were replaced with calls to a single-sided communication interface. This now allows the use of different single-sided communication libraries; for HECToR, Cray SHMEM will be used. The improved performance of the inter-task communication gives better scaling and parallel efficiency, as it is now possible to run GROMACS with millions of atoms and a lower number of particles per MPI task. GROMACS is the third most heavily used single application on HECToR, consuming 6.1% of the total machine resource and making up 4.0% of the total number of jobs.

6.7.14 Adapting QSGW to large multi-core systems

Quasi-particle self-consistent GW (QSGW) is an advanced method for *ab initio* electronic structure calculations, which is comparatively more expensive than conventional methods, e.g. local density approximation (LDA) and Dynamical Mean Field Theory (DMFT). The serial implementation of QSGW was limited to 16 atoms per unit cell, but there was scope for parallelism and the aim of this project was to develop a parallel version to enable calculations with 100 atoms. This project was proposed by Professor Mark van Schilfgaarde of the Department of Physics at King's College and was allocated 6 months support at the September 2012 panel. Martin Lueders, Leon Petit and Lucian Anton of STFC performed the work between April 2013 and September 2013.

MPI parallelism was implemented for the nested loops which calculate the central quantities in QSGW and OpenMP was used for the lower level loops. The two costliest steps in a QSGW self-consistent calculation are the calculation of the susceptibility and screened Coulomb interaction (hx0fp0), and the calculation of the self-energy contributions (hsfp0). With the new code, for the self-energy calculation (hsfp0), a 74 times total speedup relative to a sequential run of the code was demonstrated using around 1600 cores on HECToR. Also, the assembly of the screened Coulomb interaction (hx0fp0) can now be performed in 21 minutes using 2640 HECToR cores. In serial, it would take the order of days to calculate one cycle. These developments enable QSGW to perform a self-consistent calculation within several hours (or days), rather than several weeks. The new parallel LMF/QSGW (LM Suite) code will be introduced to the community through a dedicated hands-on workshop, to be held under CCP9 (Collaborative Computational Project for the Study of the Electronic Structure of Condensed Matter).

6.7.15 Parallelisation of adaptive Kinetic Monte-Carlo code

DL_AKMC is a universally applicable adaptive kinetic Monte-Carlo (aKMC) code, and is of use to a wide variety of disciplines. Recently, it has been used for simulations on radiation damage in materials. The original implementation of DL_AKMC calculated the saddle points and activation energies sequentially, however, this was restrictive for larger system sizes and longer simulation timescales (up to the order of seconds). This project was proposed by Dr John Purton of STFC and was allocated 9 months effort at the September 2012 call. David Gunn, also of STFC, performed the work between January 2013 and September 2013. The first objective was to develop a hybrid OpenMP/MPI parallel version of DL_AKMC, to enable systems with 200,000 atoms to be studied on HECToR. The second objective was to develop DL_AKMC to support a variety of potential types, and therefore increase the usefulness of the code for a wide selection of materials and systems.

The Saddle point searches were distributed using MPI, with a task farming strategy. A 2,000 atom radiation-damaged system was simulated, and a near to 100% speed up per extra core was achieved over the serial version. The Force and Energy calculations were parallelised with a combination of OpenMP and MPI. The serial and parallel calculations were compared

for 200 energy and force evaluations with a system of around 2,000 $Gd_2Ti_2O_7$ ions. Good scalability was achieved for up to 8 cores with an overall scalability of $O(N^2)$. Support for the Tersoff potential and the embedded atom method (EAM) were also added. Results were fully validated against those obtained with DL_POLY Classic. These developments to DL_AKMC will be used to conduct research funded under the EPSRC grant "Energy Materials: Computational Solutions for modelling materials of interest to the energy sector" (EP/K016288/1), May 2013-May 2018.

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

BOFFS and NEAT are two separate structured multi-block codes, which are mainly used in the design of aircraft and gas turbine engines. Both codes are used to model complex flow scenarios, and have less overhead in comparison to unstructured solvers. The first aim of the project was to implement a flexible parallel data decomposition in BOFFS to improve scalability; the second aim was to implement hybrid parallelism within the most computational parts of NEAT, i.e. the tri-diagonal matrix algorithm (TDMA), Gauss-Seidel (GS) scheme and subsidiary routines. This project was proposed by Professor Paul Tucker of the Department of Engineering at the University of Cambridge and was allocated 8 months support at the September 2012 panel. Phil Ridley of NAG performed the work on BOFFS between December 2012 and September 2013; Yiqi Qiu, also of NAG, performed the work on NEAT between March 2013 and August 2013.

A flexible data decomposition was successfully implemented in BOFFS to allow an arbitrary number of blocks to be assigned to the same MPI process. For a large eddy simulation of a subsonic jet with 50 million cells and 108 grid blocks, the run-time for ten time iterations reduced from 67.28s to 50.10s, using nearly half the original number of cores. For NEAT, a new threaded red-black ordering was implemented for the TDMA and GS routines, along with improved memory access for the main computational loops (including those in the subsidiary routines). A test case for an unsteady flow over a structured grid with 12,642,048 points and 128 blocks (i.e. 98,766 grid points per block) gave a 24% reduction in run-time. Reasonable scalability for up to 8 OpenMP threads is now achievable with NEAT in hybrid parallel mode. Furthermore, the total run-time is on average less than half that of the original code. These developments to BOFFS and NEAT will be used for the EPSRC projects EP/G069581/1, EP/I017771/1 and EP/I010440/1 (which is a collaboration between Cambridge, Warwick and Cranfield Universities).

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

ISPH is an efficient fluid solver which uses the incompressible smoothed particle hydrodynamics method to model violent free-surface flows on offshore and coastal structures. With ISPH, the simulation costs are dominated by the neighbour searching algorithm and the pressure Poisson solver. The aim of this project was to develop the ISPH software into an attractive engineering tool for complex full scale engineering problems, which are capable of handling more than 100 million particles. The work was proposed by Dr Benedict Rogers of the School of Mechanical, Aerospace and Civil Engineering at the University of Manchester. The project was allocated 9 months effort at the September 2012 call. Xiaohu Guo of STFC performed the work between December 2012 and November 2013.

The particle mapping functions were optimized for both the 2D and 3D version of ISPH by replacing MPI_GATHERV by MPI_PUT. ISPH uses a Hilbert space filling curve method with the Zoltan package to perform the domain decomposition. The next step was to optimize the neighbour list searching module by, first, updating the original data structures for the

neighbour list search kernel with a preconditioned dynamic vector approach and, second, adding a new neighbour list searching kernel with a linked list approach. This approach enables the spatial locality to be preserved, which is critical for performance. Finally, the pressure Poisson solver was improved by implementing better use of PETSc sparse matrix formats and matrix renumbering. These developments will be used on HECToR with ISPH for the project EP/J010235/1.

6.7.18 Implementing OpenIFS on HECToR

The OpenIFS project is a new initiative from the European Centre for Medium-Range Weather Forecasts (ECMWF) that will deliver a portable version of its Integrated Forecasting System (IFS) to the academic community. OpenIFS is the forecast-only part of the full IFS, which will allow UK researchers to use one of the world's leading weather forecasting models for weather and related processes, short to seasonal range forecasting, and forecast verification. The overall aim of this project was to improve the performance of the I/O operations for OpenIFS 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 support at the September 2012 panel. Mark Richardson of NAG performed the work between January 2013 and December 2013.

The original HECToR version of OpenIFS was restricted to single process I/O, making it unsuitable for production runs. Therefore, the parallel I/O capability of the ECMWF version was also implemented for the HECToR version of the OpenIFS model Fields DataBase (FDB). Furthermore, a user application was also developed to help access and management for the model output from the Meteorological Archival and Retrieval System (MARS).

6.7.19 DL_POLY_4 algorithmic improvements

This was the fifth DCSE project for improving the effectiveness of the DL_POLY_4 molecular dynamics package on HECToR. The aims of the project were to speed up the code by 10 to 30% for all users, and to reduce the time taken in calculations that involve the most expensive operations such as the long-ranged Ewald and short ranged inter-molecular interactions. The project was allocated 11 months effort at the September 2012 panel and was proposed by Dr Ilian Todorov of STFC.

To speed up performance, Ian Bush of NAG implemented an extended Verlet neighbour list (VNL) calculation within the linked cells (LC) scheme in DL_POLY_4. Asimina Maniopoulou, also of NAG implemented a symplectic multiple time stepping scheme to enable the calculation of the expensive operations (i.e. the Ewald and inter-molecular interactions) to be performed less frequently, and therefore further improve performance for certain types of problems. This work was performed between April 2013 and December 2013. The developments will be incorporated in a future release of the code. DL_POLY_4 is the fourth most heavily used single application on HECToR, consuming 5-10% of the total machine resource.

6.7.20 Further Improving NEMO In Shallow Seas (FINISS)

The NEMO (Nucleus for a European Model of the Ocean) ocean modelling code has been used successfully for a number of years in global and ocean basin applications; however its use as a shelf-sea model is less well developed. This was the third dCSE project for improving the performance of NEMO on HECToR. The aim of the work was to 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 support at the September 2012 panel. Andrew Porter also of STFC performed the work between January 2013 and May 2013. Stephen Pickles took over the work, which was completed in December 2013.

To reduce the bandwidth requirements of the 3-dimensional halo exchanges, field values beneath the seabed were removed. Performance was then benchmarked using the GYRE, AMM12 and ORCA2-LIM test cases. The facility to read in the partition information from a file was also added, which now permits the use of off-line generated “grid-partition” maps. Furthermore, the outer loops over the vertical dimension were redeveloped, so that lower levels (which are beneath the seabed) are not traversed. Finally, redundant computations on land and beneath the sea bed were removed by redeveloping the z-first ordering so that loops are only traversed over the active levels at each grid location. This also gives better load-balancing and improved cache re-use in the key routines.

6.8 Completed University DCSE Projects

6.8.1 Parallelised Uncertainty Quantification using Differential Geometric Population MCMC

The Differential Geometric Markov Chain Monte-Carlo (MCMC) sampler allows key Bayesian quantities to be accurately estimated. This project was proposed by Professor Mark Girolami, of the Department of Statistical Science at UCL and was allocated 12 months effort. Gary Macindoe, also of UCL, performed the work between August 2012 and July 2013. The Statistical Science Computing Cluster at UCL was used for the main development.

The aim was to port an original implementation of MCMC, which used the MATLAB parallel toolbox, to a scalable implementation in C. By replacing sections of MCMC with equivalent C, MPI and OpenMP code using MATLAB's MEX interface, a fully functional parallel framework was developed. The overall performance is now up to 2 orders of magnitude faster than the original MATLAB (maximum 8 core) implementation, e.g. 100,000 samples can now be generated for a Circadian model with 100 Markov chains in around 1 hour using 100 cores, whereas the previous MATLAB version would take around 2 days with 8 cores. This work will have an immediate impact on a current EU funded project (EP7-HEALTH) and several EPSRC funded projects.

6.8.2 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. This project was proposed by Professor Anthony Paxton of the Department of Physics at King’s College London and was allocated 12 months effort. Dimitar Pashov of the Department of Physics at QUB, performed the work between July 2012 and July 2013. The Dell cluster at QUB was used for the main development.

The objectives of this project were to improve the parallel implementation of TBE by implementing a parallel Hamiltonian, charge and force calculation; updating the parallel diagonalisers; and developing a CUDA interface for MAGMA to replace ScaLAPACK. Parallel implementations were developed for the electrostatic potential routines and the band structure derived quantities and a 3D Cartesian topology was implemented for the k-point decomposition. The parallel diagonalisers were updated for better use of ScaLAPACK. Finally, TBE was ported to GPUs by using the previous developments, coupled with standardised Fortran 2003 interoperability with C, for calls to MAGMA. These developments are now being used in several EPSRC-supported projects.

6.8.3 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. This project was proposed by Professor Jason Reese of the Department of Mechanical & Aerospace Engineering at the University of Strathclyde and was allocated 12 months effort. Saif Mulla, also of the University of Strathclyde, performed the work between July 2012 and July 2013. The USE-HPC and ARCHIE-WeSt clusters at Strathclyde were used for the main development.

The aim of the project was to improve fast particle dynamics for multi-scale flow simulation by developing a hybrid CPU/GPU solver, which will integrate the OpenFOAM Molecular Fluid-Dynamics simulation package and the molecular-modelling library OpenMM. Firstly, the initial OpenMM-CUDA-based hybrid solver was replaced with an equivalent OpenCL-based hybrid solver, for faster performance and compatibility with future OpenMM updates. Secondly, to reduce CPU-GPU communication, the rest of the MD algorithm was ported to the OpenMM-OpenCL platform (i.e. the OpenFOAM and OpenMM molecular representation. The final result of this project is a GPU-accelerated Molecular Fluid Dynamics simulation package with a pre- and post-processing capability, which is now being used by several EPSRC-supported projects.

6.8.4 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. This project 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, performed the work between September 2012 and May 2013. The cx1 and cx2 clusters at Imperial College were used for the main development.

The aim of the project was to improve the scalability by adding efficient parallel preconditioners and a new hybrid (MPI and multi-threaded) parallelisation. Low energy basis (block) preconditioning (LEBP) was implemented for the conjugate gradient (substructure) solvers. In addition to the LEBP p conditioning method, coarse space preconditioning (h preconditioning) was also implemented, and then combined with the LEBP via an additive Schwarz preconditioner. Furthermore, multi-threaded parallelisation was implemented via calls to the Boost library. It is now possible to achieve around a 14 times speedup for both the pressure and Helmholtz solvers. On completion of this project Nektar++ is now capable of solving a broad range of real-world engineering problems in an efficient manner.

6.8.5 Expressive and scalable finite element simulation beyond 1000 cores

DOLFIN is a finite element library which is part of the FEniCS suite of applications for solving partial differential equations. This project was proposed by Dr Garth Wells of the Department of Engineering at the University of Cambridge and was allocated 6 months effort. Chris Richardson, also of the University of Cambridge, performed the work between August 2012 and July 2013. The Darwin Cluster at Cambridge was used for the main development.

The aim of this work was to increase DOLFIN use on HPC resources by improving its scalability. The areas of development included the introduction of hybrid OpenMP/MPI parallelism; the use of HDF5 and MPI-IO for parallel I/O; and the implementation of more scalable techniques of mesh refinement. OpenMP/MPI parallelism was added via thread-safe calls to PETSc for the linear solvers. Parallel I/O was also added, with the result that I/O time is now negligible in the context of the solution runtime. Furthermore, two new interfaces were developed for 2-D and 3-D mesh parallel refinement; the wall clock time

taken to refine a 6 million cell tetrahedral mesh of a cube to 48 million cells is now around 10s with up to 2048 cores.

6.8.6 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. This project 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, performed the work between October 2012 and July 2013. The Legion cluster at UCL was used for the main development.

The aim of the project was to enhance the parallel functionality of OPTIMET for nano-device simulation and design. This was achieved by: extending code capability to 3D geometries; implementing periodic boundary conditions; and incorporating cubic optical nonlinearities for metamaterials. The original OPTIMET code was extended to enable simulations for photonic materials with 3D structures. This new application is now referred to as OPTIMET-3D. For a problem with 12 particles, excellent parallel performance and strongly scalability was demonstrated on Legion for up to 144 MPI processes.

6.8.7 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 performed the work between October 2012 and July 2013. The Theoretical Physics and Atomic and Laser Physics Groups' own cluster at Oxford was used for the development. The overall aim was to optimise and parallelise those parts of the TNT library that involve heavy computations, to enable important quantum effects in many-body systems to be studied.

Firstly, storage of the matrices was improved by implementing a more efficient usage of memory and parallelisation using OpenMP. Secondly, more efficient and scalable calculations were developed for the core functions of the TNT algorithms: the contraction and SVD operations. Further improvements were achieved by incorporating symmetry information to enable a decomposition of the tensors into sub-blocks, so that they may now be assigned to individual MPI processes. Good scalability with 8 threads can now be achieved. Based upon the performance of the original code, a 10x speedup can now be achieved for the reshape operation and a 20x speedup for the contract operation, and also for the SVD operation.

6.8.8 HPC implementation of Agent-Based Models for Cell Cultures

ABM is an agent-based implementation for the simulation of cell population growth. This project 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, performed the work between October 2012 and April 2013. The Zen Astrophysics cluster at Exeter was used for the development. The aim of the work was to enable larger populations of cells for Synthetic Biology simulations with ABM. The serial implementation of ABM already allowed the fast growing of cells within a reasonable time resolution, however, the maximum number of cells (10^8) was limited by the available memory. This limitation was overcome by implementing an MPI parallel version of ABM in order to exploit the available

memory on distributed HPC architectures and thus increase the size and complexity of the model within a reasonable time resolution.

The MPI-parallel version of ABM uses a master-slave approach: the master process handles the I/O operations and the cells on each slave process are then allowed to grow independently. The master also controls convergence by allowing the simulation to continue until either one of the processes reaches the maximum number of cells or the time limit is reached. A typical run-time for a simulation of 10^{10} cells is now about 10 minutes wall-clock time when using 48 cores and 96GB of memory on Zen. This model is 100 times larger than previously possible.

Appendix A: Terminology

TSL : Threshold Service Level

FSL : Full Service Level

SDT : Scheduled Down Time

UDT : Unscheduled Down Time

WCT : Wall Clock Time

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

SP : Service Provision

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

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

Incident Severity Levels

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

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

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

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

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

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

Appendix B: Projects on HECToR

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

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

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

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

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

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

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

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

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