N8 HPC

Annual Report 2016

Executive Summary

A survey of all Principal Investigators who have a project registered on N8 HPC was carried out over November – December 2016 which is slightly later than usual. The findings demonstrated that N8 HPC is providing a good reliable service with good support which is very much appreciated. There are some concerns over the length of time taken to run jobs and demand for the expansion of resources.

Impact of the use of N8 HPC was demonstrated through:

- Research outputs continue to be steady with 165 research papers acknowledging the use of N8 HPC resources.
- 44 grants have been submitted, prepared or awarded which specifically mention the use of N8 HPC which is an increase on 2015.
- The largest grant to mention the use of N8 HPC is an EU funded project from the University of Manchester called "Graphene Flagship" which has a value of €1 billion (£862M).
- The total value of all grants (not including the Graphene Flagship) is £28.8M an increase on 2015.
- There are 8 industry-related PhD studentships utilising N8 HPC resources, a decrease on last year.

Aims of N8 HPC

N8 HPC was funded in 2012 with 3 main aims –

- seed engagement between industry and academia around research using e-infrastructure
- develop skills in the use of e-Infrastructure across the N8 partnership
- share the asset of skills and equipment across the N8 partnership via the facilitation of networks of people

This survey has demonstrated that we are meeting these aims.

Seeding Engagement between Industry and Academia

- 19.5% of respondents stated that their projects on N8 HPC facility involve industrial collaboration with 17 individual companies named. 43% of industrial collaborators were aware of their usage of N8 HPC.
- 8 industry-related PhD studentships making use of N8 HPC resources
- Use of N8 HPC resources has been directly responsible for 3 instances of collaboration with UK industry and 2 overseas industry collaborations
- There has been 1 knowledge partnership grant which has made use of N8 HPC.

Developing skills in the use of e-infrastructure across N8

- The use of N8 HPC has enabled 15 PI's to gain access to Tier 1 resources such as ARCHER and Hartree compared to 19 in 2015.
- A further 4 PI's have applied for access to Tier 1 resources and are currently awaiting confirmation

Sharing skills across N8

- Use of N8 HPC has enabled 8 instances of collaboration between N8 universities and 7 collaborations with non-N8 UK institutions as well as 11 overseas academic collaborations.

N8 HPC Annual Survey 2016

Introduction

N8 HPC began in October 2012 with a pilot phase which lasted approximately 5 months. During this phase each institution invited selected researchers to start using the facilities with the aim of trouble shooting any initial problems and ensuring that the machine was running properly.

After the initial pilot phase, the facilities were opened up further to key researchers from all N8 institutions. The machine usage has increased to the point where the machine has been fully utilised over the last year.

In order to assess the impact the usage of N8 HPC has had on research at the N8 Institutions during the last 4 years, a survey was sent to the PI's of all registered projects on N8 HPC. The 2016 survey was again based on the previous 2014 surveys with minor modifications. The survey was again split into two parts – part one to investigate the impact N8 HPC has had on research; part two to examine how the facility was performing in terms of service.

The survey had a 50.66% response rate compared to last year's 56%. Many reminders were sent out to all PIs containing their individual survey links. The response rates from the individual institutions are shown below in Fig.1.

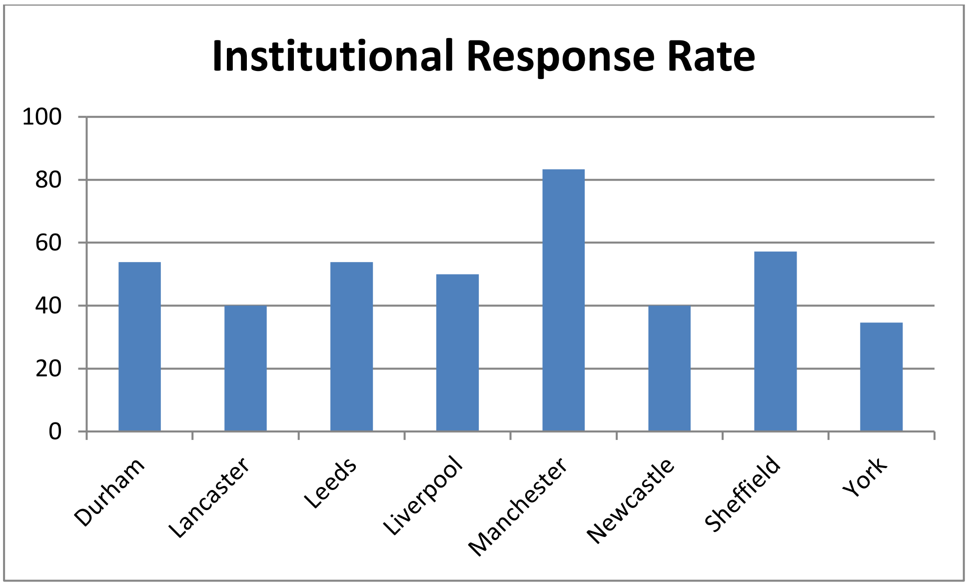

Fig. 1 Individual institutional response rates

Part 1: Impact of N8 HPC on Research Outputs

Research Papers

Q6. Please list the details of any publications (including those in in preparation, submitted, in press or published) that are associated with the use of N8 HPC.

PI's were asked for details of any papers that were associated with and / or acknowledged the use of N8 HPC.

In summary 165 papers have been produced in the last 12 months and comparisons against previous year's numbers can be seen in Fig. 2. Appendix 1 shows the number of papers produced by each N8 institution that acknowledge N8 HPC.

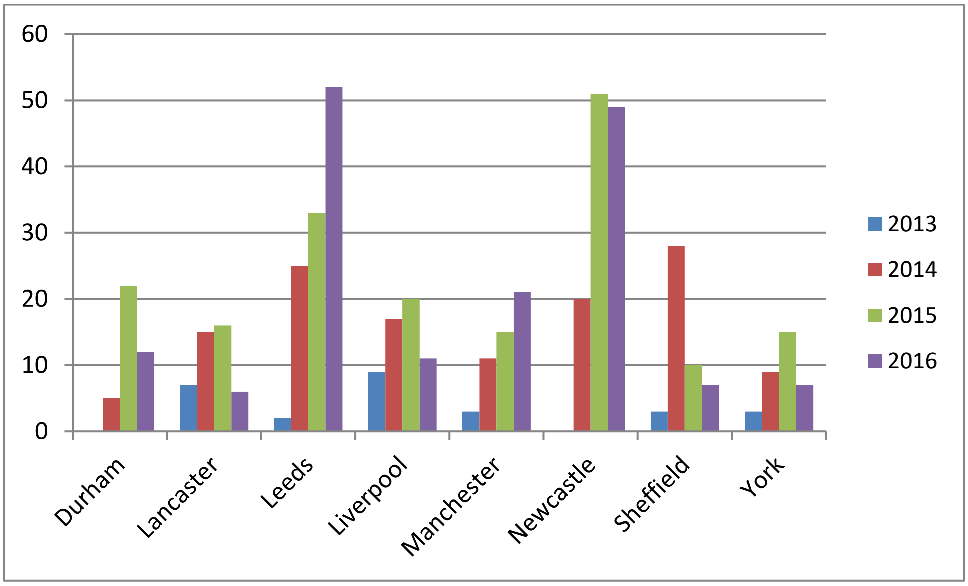

Fig. 2. Number of papers acknowledging use of N8 HPC over the last 4 years.

Research Grants

Q7. Please list the details of any grant applications which state the use of N8 HPC.

There are currently 24 active grants that state the use of N8 HPC facilities with 7 more submitted and 8 in the preparation stage. There were also another 5 grants that stated the use of N8 HPC that have come to an end in the last 12 months. The full list of grants which state the use of N8 HPC resources can be found in Appendix 2.

Industrial Studentships

Q8. Do you have any industry-related PhD studentships that use N8 HPC resources?

PI's were asked for the details of any industry related PhD studentships which made use of N8 HPC resources. 8 industry related PhD studentships were reported compared to 12 in 2015 and 13 in 2014. A full list together with the industry partners is available in Appendix 3.

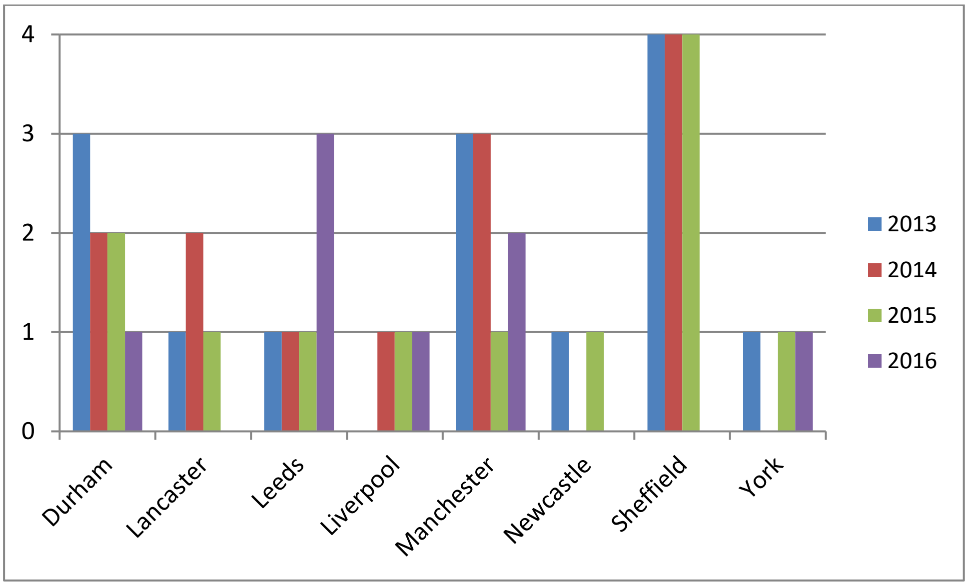

Fig. 3 Number of existing and forthcoming industry studentships at N8 institutions.

Enabling Collaboration

Q9. Has your use of N8 HPC resulted in collaborations with another researcher, institution or industrial partner?

The first three categories in the graph below refer to academic collaboration enabled by the use of N8 HPC services. There has been a slight increase compared to last year regarding collaboration within N8 universities but collaboration between academics from UK non-N8 universities has fallen. There has been a slight increase in UK industry collaboration since 2015 although overseas industry collaboration has fallen slightly.

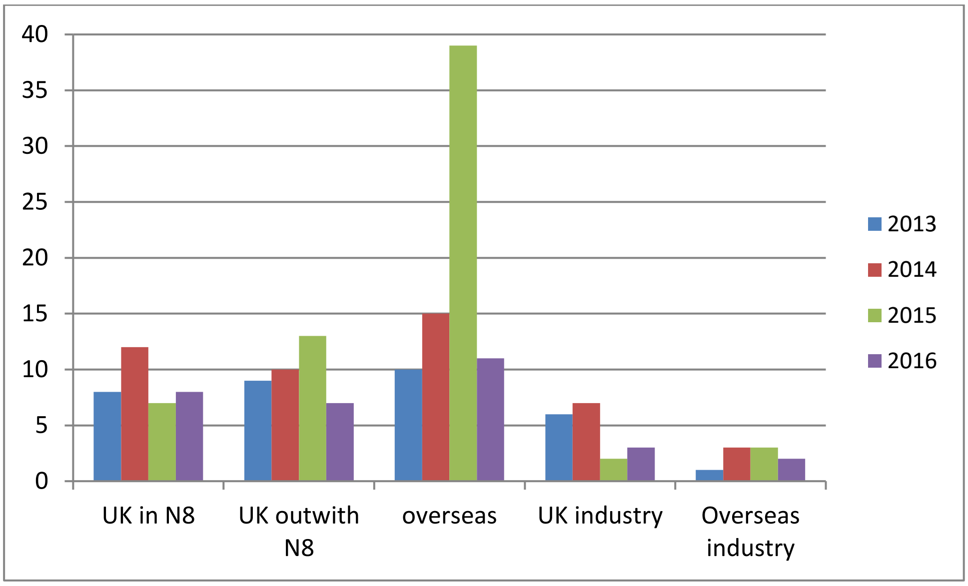

Fig. 4 Collaboration as a result of N8 HPC usage

Industrial Collaboration

Q10. Does your research using N8 HPC involve any industry collaborations (funding, joint research projects etc)?

Q12. Do any of your N8 HPC projects involve the NHS? If so please give details.

Q13. Do any of your N8 HPC projects involve the Met Office? If so please give details.

Q14. If applicable, is your industrial partner aware of your usage of N8 HPC?

In 2016 only 19.5% of respondents confirmed some sort of industry involvement compared to 24.3% in 2015 continuing the downward trend seen in previous reports. However 17 companies were named – the same number as last year.

Only one PI had any NHS involvement compared to two last year. There was an increase in Met Office involvement with 9 PIs stating they were involved with them through joint projects and studentships. This is an increase on 1 from last year.

Industrial collaborators awareness of N8 HPC usage has decreased with 42.9% of PIs stating that their industry collaborators were aware of their use of N8 HPC compared to 76% last year.

Industrial Collaborative Papers

Q15. Have you produced any joint academic-industry publications involving the use of N8 HPC? If so please provide details.

This year there are 4 published joint academic-industry publications with another 1 in preparation. This compares with 6 published and 2 in preparation in 2015.

Funding Sources

Q16. Does your N8 HPC associated research involve funding from any of the following channels? AHRC, BBSRC, EPSRC, ESRC, MRC, NERC, STFC, Charity grants, Innovate UK (TSB) grants, Knowledge Transfer Partnerships, Other (please specify).

In 2015 the largest funding category was "other". In 2016 EPSRC funding was the main source of funding but only just. In the "other" category the biggest source of funding was the EU followed by "University".

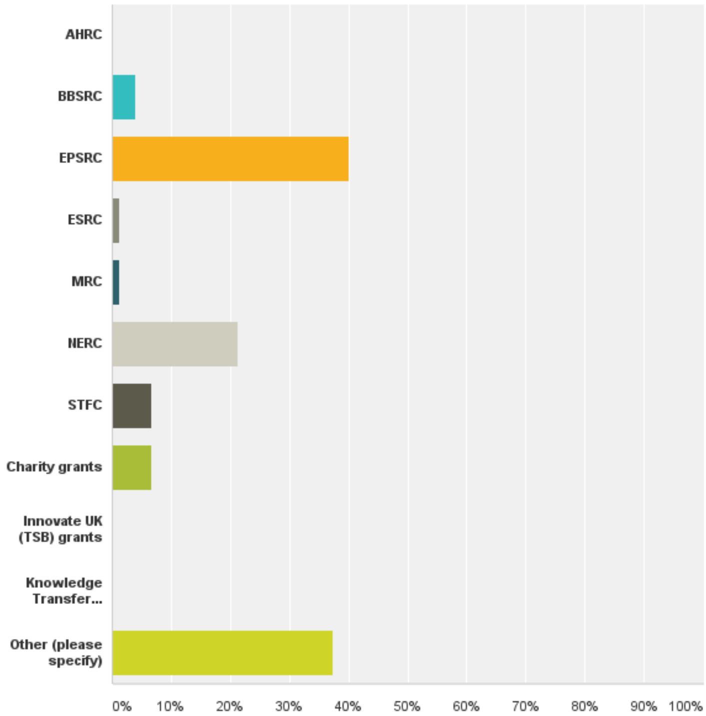

Fig. 5 Funding sources for N8 HPC projects

Knowledge Partnership Grants

Q17. Have you had any knowledge transfer projects (KTP, KTNs etc) that involve N8 HPC?

In 2016 there was 1 project that came under this category:

- Jaume Bacardit, Newcastle University
  - KTP - Applying Data Mining to create intelligent CAD tools. Industry Partner - XACT PLC. Start date: June 2013, End date, June 2015

Contact with university Business Engagement Managers

Q18. Have you had any contact from your institution's Business Engagement team regarding N8 HPC? If so please provide details. Have these discussions proved fruitful?

96.1% of project PIs have had no contact from their Business Engagement Manager down from 97% in 2015.

Interaction with CDTs

Q19. Are you involved in a Centre for Doctoral Training (CDT) at your institution? If so please give details of the CDT, your role and if N8 HPC is used or could be used within the CDT.

This was only the third year that this question has been asked. In 2015 8 project PIs identified themselves as having involvement with CDTs. This year the number was 7 with involvement mainly consisting of supervising students who are involved in CDTs.

Access to Tier 1 Resources

Q20. Has your use of N8 HPC allowed you to move upwards and gain access to any Tier 1 (ARCHER, Hartree) and Tier 0 (PRACE/Tianhe/Xsede) resources? If so please give details of the resource used.

This year 15 PIs and researchers have gained access to Tier 1 resources such as ARCHER and Hartree through their use of N8 HPC with another 4 PIs having applied for access and currently awaiting confirmation. This is in contrast to 19 PIs gaining access in 2015.

Part 2: Service Satisfaction with N8 HPC

Project Applications

Q21. How easy did you find the N8 HPC project application procedure? (1- difficult; 5 – easy)

Respondents were asked to rate the N8 HPC project application procedure on a scale of 1- difficult; 5 easy.
The average satisfaction score was 4.35 which is comparable to the previous year's results of 4.31 in 2015 showing that there is still a high level of satisfaction with the procedure. There were 7 comments with the majority being positive. There was only 1 comment that wasn't positive relating to an application that had to be resubmitted.

User Applications

Q23. If you have applied for a user account, how easy did you find the N8 HPC user application procedure? (1- difficult; 5 – easy)

Once a project has been approved it is issued with a code by the local institution N8 HPC helpdesk. This code must be entered into the user application form by researchers associated to that project who wish to apply for an account.

Respondents were asked to rate the user application procedure on a scale of 1- difficult; 5 easy. The average score was 4.53 up slightly on 4.42 from last year. There were 8 comments on the user procedure, half of which were positive about the procedure. The other comments referred to issues in obtaining user accounts for an undergrad (1) and for non-Leeds users (3).

Technical Documentation

Q25. Does the technical documentation on the N8 HPC website meet your needs?

Q26. Is there anything missing from the technical documentation or is further explanation required in any area?

The website has been present in its new form for about a year and feedback was sought on the type and quality of information provided on the website. 80% of respondents felt that the technical documentation met their needs, a decrease from 83% last year. One user felt that the technical documentation did not meet their needs but did not provide any details regarding what was missing. 8 suggestions were received as to how the technical documentation could be improved and these will be noted.

Running Jobs

Q27. Have you or your researchers run a job on the N8 HPC?

93% of the respondents or their researchers had run a job on the system compared to 86% in 2015. 5 PIs replied to say that their researchers had not run a job on N8 HPC. The respondents were spread across several sites – Durham (1), Leeds (1), Sheffield (2) and York (1).

Support and Helpdesks

Q28. Have you or your researchers required support at any point to run a job on N8 HPC?

58.7% of respondents have required help to use N8 HPC which is up from 46% in 2015.
Of the 41 people that required assistance only 1 did not receive the help they required.

*Q30. If you required support did you know how to contact the correct helpdesk (i.e. your local helpdesk)?**

The helpdesk provision at N8 HPC is devolved to the local institution with researchers asked to contact their local N8 HPC helpdesks in the first instance. The email addresses for the helpdesks are listed on the new N8 HPC website under "Help" and then under "Helpdesk". This year everyone who required help knew how to contact the correct helpdesk compared 95.5% last year.

Q31. If you have used a helpdesk in the past year, was your query handled in a satisfactory and timely manner?

Out of 41 people who contacted their local helpdesk for assistance, only 1 felt that their query was not handled in a satisfactory and timely manner.

Training

Q32. Do you or your research team require any training that could be provided locally or within N8? If so please give details.

17 PIs (8 in 2015) said that their research team would benefit from training. A list of training recommendations from the PI's is listed below –

- Fortran / MPI – 3 requests
- Parallelising code – 3 requests
- General HPC training – 3 requests
- Advanced MPI
- Code optimisation
- Unix command line
- Basic running simulations on HPC
- GPU
- Hybrid programming techniques
- Openacc

Personal Recommendations

Q33. Would you recommend the use of N8 HPC to colleagues?

98.7% of respondents would recommend the use of N8 HPC to their colleagues (a very similar percentage compared to last year). Only 1 PI said that they would be unable to recommend using the resource but no reason was given.

Looking to the Future

Q34. What is your level of readiness for making best use of the novel technologies (e.g. within the EPSRC Tier 2 & Hartree)?

This was a new questions introduced this year to gauge interest in the new EPSRC funded Tier 2 facilities. The table below shows the percentage of respondents ready for the new Tier 2 technologies.

	Ready to use now	Ready but need support around access	Will need code modification without support	Requires research software engineer effort	Will need effort from others (author, vendor etc.)	Not appropriate for our workload
GPU	14.9	9.0	3.0	17.9	17.9	37.3
Xeon Phi	10.6	9.1	3.0	21.2	13.6	42.4
ARM	9.7	1.6	1.6	11.3	12.9	62.9
OpenPower	8.1	4.8	1.6	12.9	9.7	62.9

Table 1. Percentage of respondents ready for new Tier 2 technology

Comments were also received regarding these technologies with several PIs commenting that they did not know much (or anything) about these technologies or if their code was suitable for it. A full list of comments is in Appendix 4.

Q35. Please outline the direction of travel in your research over the next 5 years and the computational requirements that go with this in order to remain internationally competitive, including compute-cycles, job size, local memory, storage and software development.

This was another new question for this year and many PIs gave very detailed answers which are contained in Appendix 5. Further analysis of the responses is required.

Q36. What activities would you like to see N8 HPC collaborate on to prepare you for the future?

There were 21 responses to this answer and the detailed responses are given in Appendix 6. The responses can be grouped into four main categories – the size of the resource, the technology, training / outreach and support.

Q37 Do you have any further feedback you would like to provide? Are there any improvements you would like to see?

11 feedback comments were received which is lower than in previous years but this years survey provided much more opportunity for feedback through other questions. Most comments could be categorised under "thank you" or "bigger machine". Full details of the comments received are in Appendix 7.

Success Stories

Q26. We are always looking for success stories and case studies to publicise on the N8 HPC website and further afield to organisations such as the Department of Business, Industry and Skills (BIS) and Innovate UK. If you have a story then please leave brief details and we will get in touch.

Several PIs volunteered to help with case studies of which there are some good leads that will be followed up by the current N8 HPC intern.

Conclusion

Overall researchers are still happy with the service but the age of the machine is beginning to tell with calls for a larger machine continuing to grow. Particular praise was given to the quality of local helpdesk support and the ease of application. However the number of grant applications citing the use of N8 HPC is decreasing with one user commenting that their local resource was now of comparable size to N8 HPC. The decrease in citations may be also be due to the lack of certainty regarding N8 HPC as the machine is now due for replacement.

Appendix 1 – list of academic output acknowledging the use of N8 HPC

Academic Papers

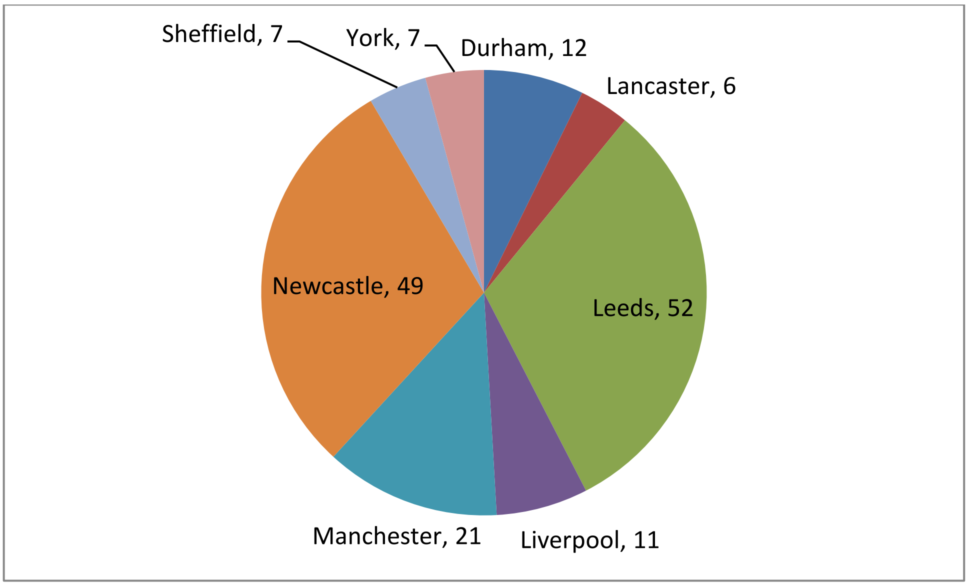

Number of papers produced by each N8 institution acknowledging N8 HPC.

Appendix 2 – List of research grants mentioning the use of N8 HPC resources

PI	Institution	Grant Title	Funding Body	Grant Status
Stewart Clark	Durham	EP/N024028/1 DFT+mu: A Step Change in Muon Spectroscopy	EPSRC	Active
Stewart Clark	Durham	EP/M010953/1 Strong correlation meets materials modelling: DMFT and GW in CASTEP	EPSRC	Active
Halim Kusumaatmaja	Durham	Wetting of Elastic Fibres: A Novel Immersed Boundary-Lattice Spring-Lattice Boltzmann Simulation Approach	EPSRC	Active
Mark Wilson	Durham	Modelling Migration of Small Molecules in and on Polymers	P&G	Active
M. Sergio Campobasso	Lancaster	Title Unknown	Archer Rap	In prep
Neil Drummond	Lancaster	Phase Behaviour and Electronic, Excitonic and Vibrational Properties of Boron-Carbon-Nitrogen Materials	EPSRC	In prep
Natalie Gilkeson	Leeds	Remodelling building design from a human centered design	EPSRC	Active
Lauren Gregoire	Leeds	Searching for the deglaciation: spatio-temporal boundary condition uncertainty and its implications for understanding abrupt climate change	EPSRC	Active
Zlatko Papic	Leeds	Strongly-Entangled Topological Matter	EPSRC	Active

Wuhu Feng	Leeds	Cosmic Dust in the Terrestrial Atmosphere (CODITA)	ERC	Active
Andy Nowacki	Leeds	The Secret History of the Earth's Mantle	Leverhulme Trust	Active
Ruza Ivanovic	Leeds	Independent Research Fellowship #NE/K008536/1	N/A	Active
Ruza Ivanovic	Leeds	Forward modelling of past abrupt climate transitions.	NERC	Active
Martyn Chipperfield	Leeds	The Global Methane Budget	NERC	Active
Steve Arnold	Leeds	Air Pollution and Arctic Climate: Improved Process Understanding for Prediction and Mitigation (AirPAC)	NERC	Submitted
Ming Li	Liverpool	Sand Transport under Random and Breaking Wave Conditions	EPSRC	Ended
Gilberto Teobaldi	Liverpool	INSPIRE Physical Sciences: A Synergy for Next Generation Materials Science.	EPSRC	Ended
Ming Li	Liverpool	Towards Realistic Modelling of Sand-Mud Mixture Sediment Transport	EPSRC	In Prep
Ming Li	Liverpool	Towards Realistic Modelling of Sand-Mud Mixture Natural Environment Sediment Transport	NERC	In Prep
Alistair Revell	Manchester	Hybrid RANS-LES for Complex Geometries	CD-adapco	Active
Viktor Zólyomi	Manchester	Centre for Doctoral Training in the Science and Applications of Graphene and Related Nanomaterials (GrapheneNOWNANO)	EPSRC	Active
Viktor Zólyomi	Manchester	Engineering Van der Waals Heterostructures: From Atomic Level Layer-by-Layer Assembly to Printable Innovative Devices	EPSRC	Active
Viktor Zólyomi	Manchester	Graphene Flagship	EU	Active
Viktor Zólyomi	Manchester	Designer Nanomaterials Assembled from Individual Atomic Planes.	Lloyds Register Foundation	Active
Alistair Revell	Manchester	Development and Implementation of Embedded LES into OpenFOAM	National Grid	Active
Alistair Revell	Manchester	PELskin: a novel aerodynamic coating	Pelskin EY	Completed
Peter D. Lee	Manchester	EXOMET: Physical Processing of Molten Light Metals Nano-Particulate Reinforced Alloys under the influence of External Fields	EU	Ended
Bill Sellers	Manchester	Returning to the Ocean: The Transitional Mechanics of Legged Whales	Leverhulme Trust	Under Review
Bill Sellers	Manchester	Inside out: unravelling the Miocene Primate Locomotion through their internal skeletal design	Marie Curie	Under Review

David Schultz	Manchester	Palegeography and Palaeoclimate Through The Ages: The Relative Importance of Continental Configuration, Solar Constant, and Atmospheric Composition	NERC	Under Review
Graeme Sarson	Newcastle	Wetting of Elastic Fibres: A Novel Immersed Boundary-Lattice Spring-Lattice Boltzmann Simulation Approach	EPSRC	Awarded
Paul Bushby	Newcastle	The Research Programme in Astronomy Theory	STFC	Awarded
Nilanjan Chakraborty	Newcastle	A combined experimental and numerical investigation into the natural convection of viscoplastic fluids	EPSRC	In Prep
Nick Proukakis	Newcastle	Bose-Fermi Ultracold Mixtures	EPSRC	In Prep
Carlo F Barenghi	Newcastle	Quantum Vortex Reconnections	EPSRC	Submitted
Nilanjan Chakraborty	Newcastle	Adaptive software for high-fidelity simulations of multi-phase turbulent reacting flows	EPSRC	Under Review
John Harding	Sheffield	NE/M011429/1	NERC	Active
Ning Qin	Sheffield	H2020 DRAGY	EU	Awarded
John Harding	Sheffield	EEP/I001514/1	EPSRC	Ended
Shuisheng He	Sheffield	European Small Modular Supercritical Water Reactor Technology	EU	Submitted
Keith McKenna	York	Optimisation of Charge Carrier Mobility in Nanoporous Metal Oxide Films	EPSRC	Awarded
Keith McKenna	York	Non-Equilibrium Electron-ion Dynamics in Thin Metal-Oxide Films	EPSRC	Awarded
Jon Hill	York	Bioshield	EPSRC	In Prep
S A Cavill	York	Straintronics	EPSRC	In Prep

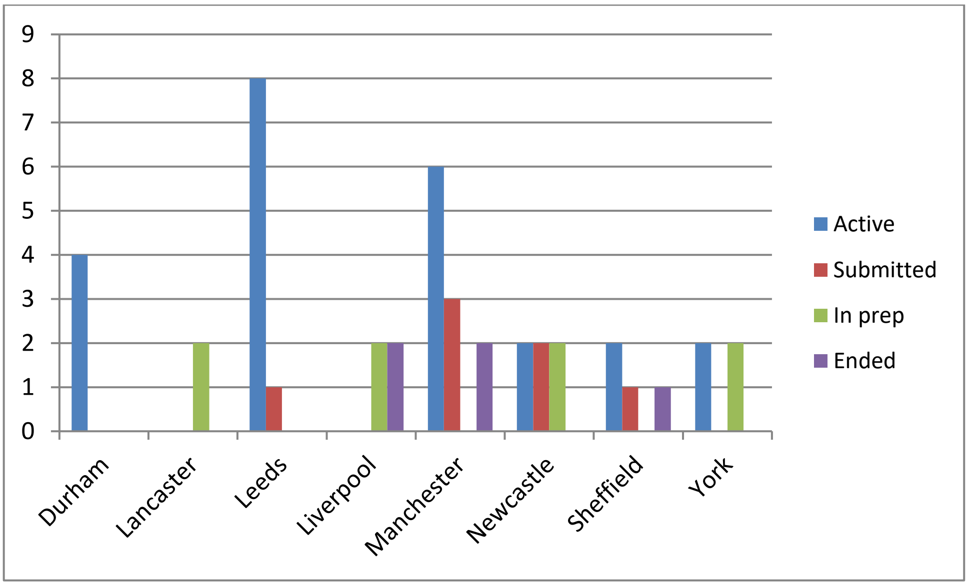

Fig. 5 Number of research grants and their state at the time of the survey per institution.

Appendix 3 – list of industrial studentships involving N8 HPC

PI	Institution	Funders	Industrial Involvement	Title	Dates
Mark Wilson	Durham		P&G	Simulation studies of polymer-surfactant phase diagrams	1st Oct 2014 - 31st Dec 2017
Alistair Revell	Manchester	iCASE	National Grid	Application of recently developed Embedded Large Eddy Simulation for detailed investigation of flow physics	2014 - 2018
Alistair Revell	Manchester		McLaren F1		
Jaafar Elmirghani	Leeds	EPSRC CASE	AstraZeneca	Molecular simulation of crystal growth & dissolution	July 2014-Dec 2018
Jaafar Elmirghani	Leeds	Syngenta	Syngenta	Solubility prediction from molecular simulation	Jan 2016-July 2019
Jaafar Elmirghani	Leeds	Unilever	Unilever	Towards modulating the skin barrier: mechanistic insights into tight junctions using computer modelling	Oct 2015-Sept 2018
Ming Li	Liverpool	NERC/Met Office CASE studentship	Met Office	Solar and space weather impacts on atmospheric chemistry and connections to surface weather and climate.	1 Oct - 2016 30 Sept 2020
Emma McLarnon	York	BBSRC	DLF Trifolium	Phenotypic and molecular characterization of silicon uptake and deposition in forage grasses	Oct 2013 – Sept 2017

Appendix 4 – preparedness for Tier 2

I do not have much experience using GPUs, but one of the codes I use a lot can make use of Nvidia GPUs. It seems that many codes have been locked into one vendor and I would like to see a move to OpenCL, and support from HPC vendors should be given for this (for their own benefit in being able to choose other GPU suppliers).

In my group we have started thinking about GPU and XeonPhi and made some exploratory ports. We have no experience of ARM or OpenPower and do not currently have access to any hardware and have no knowledge of the relative merits of this approach or its weaknesses. Need a lot more info here.

My software should run on Xeon Phi but current access models seem to make it relatively tricky to actually get code to compile and run as easily as it should. I don't know whether it would be possible to implement some sort of back-filling for high throughput applications that can run when cores are not being fully utilised.

Our code is very portable but I have not looked into this different hardware.

Our software supports Intel CPU

Probable that we could use OpenPower but require further details

Will XeonPhi be available on Polaris? Will GPGPUs be available? Generally the codes in this domain (earth system modelling) are Hybrid MPI and OpenMP with majority being MPI only. Any novel architectures will require significant effort to port codes although there are efforts elsewhere (USA for WRF and UK Met Office for UM) but are at differing levels of maturity.

x86 is necessary. Novel technologies still a toy.

I do not know anything about these technologies.

I don't know what most of these features are and whether my code would benefit from them.

Not familiar with Xeon Phi, ARM and OpenPower and their capabilities

Appendix 5 – Direction of travel in your research over the next 5 years

>256 core jobs with at least 4GB/core RAM

A mix of job types: Small: 20-60 cores, 2-6 hours, 1 GB/core Medium: 60-500, 6-24 hours, 1-2 GB/core Large: 500-2000, 6-24 hours, 1-2 GB/core High throughput is just as important as capability for large jobs in order to increase productivity.

Adaptive mesh refinement (AMR), multi-physics multi-scale fluid flow problem (e.g. droplet combustion with atomisation) with higher Reynolds number: This translates to more cores and AMR needs serious software development

CPUs usage ~256 cores per job, 10 jobs at any given time; Current local memory is fine. Storage is becoming critical as trajectories generate large amounts of data - requirement 100s TB.

Dynamics far from equilibrium for multicomponent quantum mixtures Memory: ~10 GB Storage: ~few TB Options to subcontract dedicated N8 software engineers would be very useful

Exploring new ways to address some current limitations to computational performance limiting ability to capture potentially important "multi-scale" physics. These simulations are currently not routinely feasible -- will be exploring both computational techniques (improving software, novel algorithms, new hardware) as well as employing physics insight to simplify. Will require ongoing access to current levels of resource, with the ability to occasionally access substantially increased resource for short periods (order of magnitude increases but for ~ 1-2 months).

For our moderate use purposes, and in my limited experience of using Polaris so far, no significant improvements to the service are needed. It should be ensured that demand can be matched to supply at the current levels.

For the particular project this will require very large memory nodes which are massively parallel. For other research having a large number of well connected nodes is useful.

Good turnaround of medium size jobs - up to 500 cores.

Hi Resolution Earth simulators require hi number of nodes and memory per core. They often run an ocean model alongside an atmosphere model. High speed data transfer for large files to other establishments for post processing. Regional air quality requires complex chemistry so flop rates are significant

I am only a PhD student so I am not sure what I will need over the next 5 years. I believe storage will need to increase as we continue running ever longer jobs. I anticipate needing more HPC time for jobs on medium number of cores (32-100) with access to machines that are large enough to facilitate many of these runs with a queuing system that allows multi core jobs to get on the machine easily.

I anticipate modest growth and a focus on improving our existing software without incurring large additional costs. Typical job-size 128 to 256 nodes.

I may need to run the same number of simulations to collect data on another type of instability I am researching which will produce around 1Tb of data and several weeks of computing time in total

I will need the ability to run long experiments (over weeks/months) without spending most of time queuing.

In computational chemistry there is a trend towards larger systems with polarisable forcefields. Computer requirements will grow with this.

In the next few years, I envisage my own research largely continuing to make use of Tier 2 resources, increasing relatively slightly compared to current resources. (Our applications make use of novel analysis, more than requiring the greatest computing power.) For some specific applications, however, I do plan to move towards Tier 1 resources.

Jobs perhaps using up to 128 to 512 CPUs but still running for substantial wall clock time.

Jobs up to 128 cores. Total memory 64Mb in some cases.

Job-size is critical - more processors means higher resolution/larger computational domains/more realistic parameters.

Large memory for core.

Large memory nodes (>=512gb RAM) with longer run time (above 48hrs)

large memory, storage, fast speed

Long run times with large local memory, CPU-bound applications

Long runs of many months-years walk-clock time of parallel jobs requiring 16-100 processors each. Local memory usage is low. Multi-month long (wall-clock) parallel jobs of ~1000 processors. Again, local memory usage is low.

Memory requirements are expected to increase. More physical memory per compute node will be required to remain competitive. For parallel jobs, faster interconnects between nodes will be needed to handle high-end calculations.

Moores law: double every 18 months!

More nodes would be required (we are limited by our licence for a commercial code)

Most of my research involves very large simulations: typically about 5-15k cores running for ~2hours. Memory requirements are usually only ~100MB per core.

Moving towards real-time simulation using GPU. Requirements of 16 cores per user with a minimum of 2GB per core.

Much of our research is based on CFD: Direct numerical simulation (DNS) and large eddy simulation (LES), which require very high computing power. Our research is therefore very much dependent the availability of HPC resources.

Multiple longer MD simulations require trajectory files that can add up to terabytes of disk. Mostly GPU powered. However our code relies on fast interconnect CPU currently.

My plan is to run a small number (i.e. 1 per month) of simulations at high resolution. Each simulation requires about 0.4M cpu hours and about 20Gb of storage.

My research is in data analytics and big data, so computational resources are totally crucial for my work, and my demand will only grow. What would be useful to me is to be able to use the N8 resources via big data frameworks such as Hadoop or Spark.

My research on materials involves simulations done on different grid sizes. Currently I run simulations on 50 x 50 x 50 grid points in 3D, which corresponds to a very small grain size in reality. So I have to scale down my 3D simulations to 2D in order to simulate large grids (large grains) with less queuing time. Currently, I have tested a maximum of 1000 x 1000 grid points in 2D, which also increased the queuing time considerably. In the future, I would like to simulate large 3D grids with less queuing time. So it would be beneficial to increase the number of nodes and number of CPU's per node and reduce the queuing time.

my research will be focused toward novel PV concepts and novel QIP device. I expect my typical jobs in next 5 years to be 512 CPU 10-100 GB ram using my MPI parallel kph code

N8 access required only for a few more months on this project code.

Need flexible compute and storage to scale from very small to very large

Neural networks keep gaining wider adoption in language processing. The major Neural Network packages should be installed as standard, e.g., TensorFlow, Torch or DyNet (formerly CNN), while the use of GPU is now considered to be standard practice.

NGS analysis of whole genome resequence data. Low RAM requirements, 4gb. Storage requirements, 1Tb. High cpu requirement, several 100 cpu.

Our models will increase core number from 10s to 100s and will require large RAM available (16GB per core typically)

Our need for high performance computing will grow to allow use model more complex networks

Our research will be largely focused on developing Computer Aided Materials Science (CAMS). We envisage CAMS will benefit from (1) state-of-the-art computational techniques (e.g., GPU-accelerated quantum chemistry calculations), (2) a steady increase in the computational resources (e.g., number of CPUs, tolerance to flush storage of large volumes of files and large amount of data in the working directory) available to us, and (3) highly efficient data I/O and transfer.

Part of my research will continue with Renewable Energy Fluid Machinery Engineering supported by HPC_enabled CFD. The wind energy part of this work will require further code development work, verification and validation of which require simulations ideally suited for a tier 2 HPC resource like Polaris. This resource is also important for preliminary medium-size engineering analyses and parallel performance verification/testing before applying for ARCHER access.

Predominantly chemistry of complex inorganic materials. Our use of ~100 millions AUs of Archer time per 6 months is realistic for the next 2 years, but we also use as much of local machines and Tier2 as possible, jobs can be anything from 64 cores to thousands depending on complexity of materials and composition. Storage is not an issue, but RAM will become a problem as the system sizes grow. Not sure we can put a figure on it, the bigger the computer the more complex the system we'd consider.

Probably going to remain the same, in the next few years.

Processing power in terms of number of cores Fast MPI

Range of job size according to need. Making use of novel architectures.

Remains the same

Research is likely to expand considerably when I leave my current post as head of department, jobs sizes are likely to be ~128 cores, 1 Gbt per core, and store typically ~ 1 Tbt of data; systems will require run lengths of ~month on today's resources (128 cores)

Software development needed. Less queue time

Software development primarily for enhancing code capabilities: novel meshing strategies, and improved algorithms for timestepping. Low-level optimisation for new hardware generally a lower priority than getting the algorithms working to start with,

The computational requirements for electronic structure calculations are likely to remain high if one wants to get results within a reasonable time frame. I envision 2400-core hour jobs becoming routine with 3-4 GB of memory per core and up to 100 GB of storage being required per job.   As for future software development, the codes are currently optimised to run efficiently on multi-core machines via MPI. Significant software development and testing would be needed to change paradigm.

The computing needs in my group vary depending on the research problems. In the short term, I anticipate that my group will continue to run C++ MPI jobs (20 jobs per month for 3-4 days) of 32-64 cores per job. Each simulation will need about 1GB of memory and 3GB of storage. We also run a large number of single CPU Monte Carlo simulations, but they are probably less appropriate for Tier 2 facilities. In the medium term, we are looking to develop a GPU version for our lattice Boltzmann code.

The main concern is the number of cores available, both for running large jobs and large numbers of independent jobs. In 2-3 years, we will need full time access to 200+ cores.

The main emphasis will be on electronic structure calculations and the compute resources need to follow the international development

The project has finished, and the post-doctoral researcher in charge of this analysis has left. While this project was highly successful, we do not anticipate any other use of HPC platforms in the near future.

There is an increasing need to run larger jobs. Over the next 5 years I anticipate that jobs of 1000 cores will become much more common , if not the norm for HPC

To remain competitive, we need a flexible mixture of computers at school level, university level and N8 level

To remain internationally competitive in my field, tens of millions of core hours per year are required. Typical job sizes will involve hundreds or thousands of cores. A typical project will require short-term storage requirements of a few TB and long-term storage requirements of tens of MB.

Upwards? We can always use more clock cycles and RAM than are available in a Tier 2. We have collaborators using our code (CASTEP) with 10,000 - 20,000 cores. We are looking to develop new parallelism (e.g. GPU and Phi) and extend the recently added OpenMP hybrid layer (on top of the long-standing MPI layer)

Utilization of N8 by research student is very beneficial, but level of computational requirement is relatively minor

We are expecting to scale up the currently ~million particle based simulation towards ~billions particles for more realistic condition. That means the job size can be 10-100 time larger and storage of data will be ~Tb per run.

We are growing in both model complexity, but the largest growth is in dataset size of our experiments and the need methods to handle this (20-50Tb/expt)

We are increasingly running perturbed physics ensemble simulations, as GLOMAP is quite a simple model, this means we are running a large number of small jobs. On ARCHER we have looked into grouping lots of small jobs together to make queuing easier. But we would like to retain the flexibility of running a smaller (less expensive) model as this is our testbed for developing the more complex (and more expensive) model.

We are running the WRF-Chem atmospheric chemistry transport model on 64 cores. Our principal constraints are run time (model runs exceed 48 hours, so we need to break jobs up) and I/O (storage requirements are large). We intend to increase core number in future, but do not require substantially larger resources beyond this.

We are still in the testing phase of the code and it is difficult to forecast our computing needs over the next five year. Our computational requirements for now are (per simulation): Virtual Memory: 512-1024 GB Processes: 64-128 Wall clock time: 24-48 hrs Disk memory: 2-6 G We run 3 to 5 jobs a week and will be running up 10 jobs a week after the testing phase.

We are writing proposals to study paleoclimates. If funded, we would like to significantly expand our use of N8 resources. In order for us to use N8 for the projected research, we would need more processors (32 processors to 320 processors) and a queue for jobs with a maximum job length of greater than 48 hours.

We plan to become more bio-realistic and this will inevitably increase our computational demands. We are also looking at multiphysics simulations.

We will plan to make use of Polaris resource in using three main models: CESM and WRF-chem (existing project) and the Met Office UM (will be new project). We expect 2-4 extra users of the former 2 models over the coming year.

Will require vanilla parallel processing but more of it. Not possible to give detailed appraisal over 5 years but will involve (1) simulation of structures of complex ceramics especially interfaces, (2) crystallisation (very long timescale simulations, advanced sampling techniques) (3) other long timescale problems such as diffusion in complex materials

Appendix 6

Size of the Resource

- More compute power, please.
- Closing the gap with HPC resources available at competing institutions/University network
- To provide a regional HPC with medium range but a substantially improved system in its scale.
- Polaris is very old. It was too small to serve all of N8 when it was new. That is even more so now. It needs to be replaced with something at least 10x bigger.

Technology

- GPUs
- Mixed mode medium sized HPC for capacity (esp. GPU clusters)
- GPU provision. Visualization techniques.
- The next generation of HPC technique, such as GPU, OpenPower in terms of beneficial to the simulation efficiency and accuracy etc.
- Open to alternative models or cluster management beyond the classic SGE.
- More available GPU resource also

Training / outreach

- HPC training for new researchers
- Training and code modification support for new architecture.
- University-level training of Math/PhD students and postdocs on the use of parallel machines
- Training activities
- The N8 networking events in York are very useful.
- For N8 HPC to run regular workshops on thematic research areas
- Develop centre of excellence to enable scientific collaborations.

Support

- It might be useful to be able to hire software engineers on grants. Perhaps this is already possible.
- eCSE-like support for new HPC systems such as GPUs.
- Software development of existing codes to better utilize GPUs.
- Would be nice to see more bespoke knowledge of different versions of the Met Office Unified model including ability to port new and older versions of the model onto new N8 hardware.

- Installation of auxiliary software   Creation of project programming environments   i.e. module load PrgEnvUMvn8.4   or similar. Perhaps Virtual machine creation?

Miscellaneous

- Industrial access is a good angle.

Appendix 7 – feedback comments received

Thank you

- I am very happy with Polaris. It is a lot easier to get jobs running than it is on Archer; the queues are smaller, and the machine seems to have fewer failures.
- I really like the service.
- I've always been very happy with the N8 service. I believe it plays an important role in the research HPC ecosystem in my field (computational fluid dynamics)
- No comments, but just to say thank-you and the service is appreciated.

Wall clock

- I struggle with the 48-hour wall clock limit. 48 hours is often not long enough to complete enough of one of my jobs to make analysis possible, which makes it difficult to manage my time (planning ahead) if the job will have to be resubmitted to continue multiple times to be able to have enough data to begin analysing the output. I do not need the job to be finished, always, to start analysis, but I need a certain amount of it to be completed. The 48-hour wall-clock limit is too short for this when the machine is busier.
- Time-limitations per job (2 days) can be quite restrictive, and would be most useful to have a check-pointing software (DMTCP) available.

Software

- Installation of OpenFOAM 3 and working version of SWAK4FOAM.

Bigger machine

- More compute power, please.
- Please increase the number of nodes, CPU's per node and reduce the queuing time.
- Refresh the machine.
- There is less and less need to run our models on N8 as our Tier 0 is close to the size of the N8