



# University of HUDDERSFIELD

## University of Huddersfield Repository

James, Yvonne, Holmes, Violeta and Munnings, Daniel

Campus HPC Network Design and Monitoring

### Original Citation

James, Yvonne, Holmes, Violeta and Munnings, Daniel (2013) Campus HPC Network Design and Monitoring. In: The 15th IEEE International Conference on High Performance Computing and Communications (HPCC 2013), 13-15, November 2013, Zhangjiajie, China.

This version is available at <http://eprints.hud.ac.uk/id/eprint/21058/>

The University Repository is a digital collection of the research output of the University, available on Open Access. Copyright and Moral Rights for the items on this site are retained by the individual author and/or other copyright owners. Users may access full items free of charge; copies of full text items generally can be reproduced, displayed or performed and given to third parties in any format or medium for personal research or study, educational or not-for-profit purposes without prior permission or charge, provided:

- The authors, title and full bibliographic details is credited in any copy;
- A hyperlink and/or URL is included for the original metadata page; and
- The content is not changed in any way.

For more information, including our policy and submission procedure, please contact the Repository Team at: [E.mailbox@hud.ac.uk](mailto:E.mailbox@hud.ac.uk).

<http://eprints.hud.ac.uk/>

# Campus HPC Network Design and Monitoring

Yvonne James  
University of Huddersfield  
Queensgate  
Huddersfield, England  
y.james@hud.ac.uk

Dr. Violeta Holmes  
University of Huddersfield  
Queensgate  
Huddersfield, England  
v.holmes@hud.ac.uk

Daniel Munnings  
University of Huddersfield  
Queensgate  
Huddersfield, England  
d.munnings@hud.ac.uk

**Abstract**— The needs of the research communities in research institutes and Higher Education (HE) establishments are demanding evermore powerful computing resources for supporting complex scientific and industrial simulation and modeling, manipulating and storage of large quantities of data [6,9].

In this paper we present our experience at the University of Huddersfield (UoH), UK in developing the HPC systems infrastructure, removing a technical burden from researchers and enabling quicker and more insightful research outcomes. We have designed and implemented the University of Huddersfield, Queensgate Grid (QGG) campus grid [7].

In the process of building QGG systems and optimising its performance, we have designed and implemented a reliable network system infrastructure. The network topology was re-designed in various stages of system deployment resulting in a reduction of the number of switches, routers and network interconnects. This has led to an improvement in data transmission, a reduction in the possibility of bottlenecks and much reduced data loss [2, 9].

The rapid expansion of our campus grid has led us to question the energy efficiency of our HPC systems. Our initial investigation has targeted the transfer of data and power usage with a view to extending this work to incorporate other metrics, which is the subject of further work.

*Keywords* – HPC network design; performance; topology; energy efficiency; green computing

## I. INTRODUCTION

To respond to scientists and researchers demands, the high performance computing (HPC) resources, in form of cluster, grid and cloud computing systems are required by the institutes. The HPC resources deployment and management present challenges in an effort to achieve optimum utilisation, and deal with an ever increasing energy consumption needed to power and cool the HPC equipment.

Researchers at the University of Huddersfield currently have direct access to more computing power than ever before. The

HPC facilities are closely tailored to the needs of users – an important factor in attracting high-calibre researchers to the university. Calculations that would have taken weeks or months on a desktop machine can now be carried out in hours. Cutting-edge research in fields such as molecular biology, accelerator physics, engineering fluid dynamics, computational chemistry, image rendering and informatics are reaping the benefits.

This computing power was provided by establishing the University of Huddersfield Queensgate Grid (QGG), as seen in figure 1, which enables access to local and national HPC services.

- Local - Campus grid Queensgate Grid compute clusters and Condor pool
- National - A share in an IBM iDataPlex system as part of the STFC enCore cloud service at Daresbury [11] Laboratories, Hartree centre – IBM Blue Gene [12]
- National
  - The National e-Infrastructure Service and UK-NGI
  - The North West Grid

As a result of this development, research across the University of Huddersfield has increased in terms of the number of users accessing the HPC cluster systems, and their research output.

While these clusters were able to provide the much needed resources to an ever expanding research community, they also raised questions about their green credentials especially in terms of power consumption of our local HPC resources.

HPC systems have benefitted from the emergence of power saving and power reduction devices, which are able to hibernate when not in use and therefore save on power consumption [3]. The metrics employed to calculate performance are investigated in isolation, so the cumulative effect is missed. The Green Grid promote the use of power usage effectiveness (PUE), and data centre compute efficiency (DCcE) [2] provide metrics that are key to understanding the effectiveness of systems housed in a data centre. The need to measure these metrics has arisen from the need for HPC systems to become more affordable while still maintaining scalability and availability [5,9]. Most HPC benchmarks consider the effectiveness of node performance, CPU and

memory utilisation, but do not consider how much data is transferred in relation to power usage within the HPC network.

centres, already house the university computer services servers and provide the necessary cooling and power, and are therefore

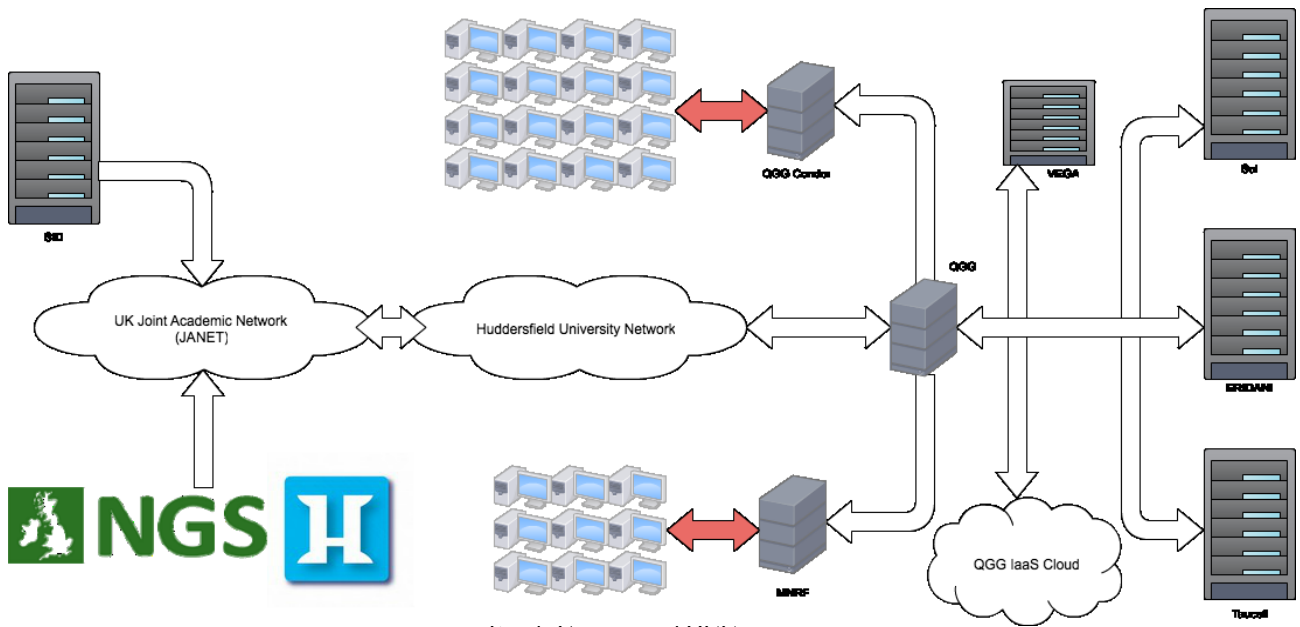


Fig. 1. Campus grid QGG

In this paper we will consider the impact of a new cluster into any HPC campus network and the implication on the network design and data transfer between various components of the HPC network. We are presenting our experience of integrating the Sol cluster into the QGG at Huddersfield.

While this paper does not specifically deal with the green credentials of an HPC system, it does consider the relationship between data throughput and energy consumption.

## II. CAMPUS HPC INFRASTRUCTURE

When designing HE campus HPC systems often insufficient consideration is given to the placement of equipment, which tend to be localised within the individual departments without provision for future expansion. As a result most of the departmental HPC clusters have slow network links within the campus network which result in bottlenecks and data loss.

An increase in bandwidth would help to improve this; however, significant improvement could be achieved by the allocation of more space within the university data centres. This would enable and provide a significant increase in bandwidth, and establish a direct connection to the network backbone.

Based on our experience, the most efficient approach in providing the HPC infrastructure for HE institutions is to move away from the “mini-data centres” located in the individual departments. Often departmental clusters are purchased from the research groups funding, but do not provide a centralised HPC resources for the benefit of the entire institution. Centralising the HPC systems would avoid the duplication of hardware and software resources, and reduce the cost of cooling and power infrastructure. The university’s data

the most suitable locations for the HPC resources.

Based on this rationale, the latest addition to our HPC resources, Sun cluster Sol, was placed in one of our data centres. We will focus on the integration of this cluster into our QGG infrastructure and the implications this would have on the campus network.

### A. Initial Campus HPC Architecture

The installation of an HPC system grew out of the need of various researchers pursuing work which requires serial and parallel processing to handle complex instruction and data sets. Various servers had been purchased to create small clusters within different departments at the University of Huddersfield. Early research [7] examined the impact of these small clusters. Our current research builds on the results of this work.

Further funding was acquired to bring the HPC resources together in one place as well to provide a further larger cluster of 158 cores; each node having 8Gb RAM and Intel Core 2 Quad 2.331GHz Processor. [7]

The latest addition to the HPC resources is a 256 core AMD Sun Systems Sunfire X4100 cluster. The installation of this new cluster led to considerations for network topology changes. Initial plan to install Sun cluster in a small data centre in the School of Computing and Engineering would have put an extreme load on the power supply and posed some security issues in terms of accessibility to the physical machines. Subsequent deliberations were driven by the analysis of the existing system and identification of the system flaws. This in turn led to a decision to allocate a space in the university data centre. This solution provided increased security, necessary power and climate control environment, and direct access to the university network backbone. The decision to house the

new cluster in an existing data centre was predominantly a management decision. However, this presented issues in terms of providing the necessary power and space to accommodate a new cluster.

### B. Existing System Analysis

The issues identified in the existing system were related to the transfer of large data files and highlighted the constraints of a limited infrastructure. The original HPC network in the School of Computing and Engineering provided a 100Mbps link to the first available switch. This then connected to a 1Gb link, and fed into the network backbone which operates at 10Gbs.

The first concern with placing a new cluster within this topology was the 100Mb links which have the potential to cause bottlenecks, performance degradation and data loss as shown in Figure 1. The second concern was that the authentication server was located elsewhere in the same building, so initially data was being pushed along a 100Mb link to this server before being sent to the new cluster traversing a number of switches in the process.




Fig. 2. Simplified existing system setup showing potential 100Mb links for bottlenecks

A further concern was the difficulty with transfer of large data files across the campus network. Due to the volume of daily network traffic, together with the additional HPC traffic, the network suffered a drop in performance, which resulted in some packets being lost. This in turn caused TCP to resend data. Network analysis was conducted using Solar Winds and Wireshark [13] software to identify that 100Mb links were at the heart of the problem. These results helped to strengthen the technical case for high speed data links.

Much of the existing HPC systems had been installed in a small data centre in School of Computing and Engineering which was connected with a 1Gb link back through several other switches until reaching the backbone. This gave rise to issues of latency especially given that a new cluster installation was expected to handle large file sizes. This issue alone was sufficient to prompt considerations of alternative locations. Figure 2 shows the initial design layout and the initial connectivity back into the main network.

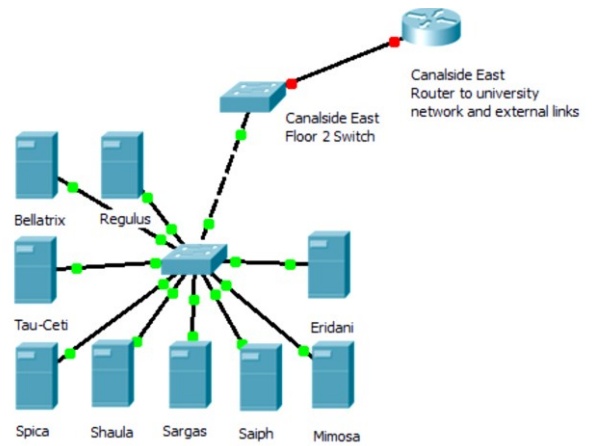


Fig. 3. Initial design for first HPC installations within the HPC research office

This original design comprised a number of virtual machines and network area storage alongside two existing clusters. As a result of the issues identified it was decided to place the new cluster into the university data centre. The new HPC network design proposed the creation of a separate subnet. This subnet benefited from a 1Gb link to the network backbone which operates at 10Gb. and provided connectivity across the campus network as well as external links.

### III. NEW HPC INSTALLATION

The university data centre now houses two racks of Sun Systems Sunfire X4100 amounting to 256 cores. In addition there are two 16Tb network area storage servers (NAS), one in the data centre and one in the HPC research office. A Cent OS server which manages authentication and some routing for other external grid services and layer three switches which form the data centre infrastructure. One switch provides the direct connection to the campus network as shown in figure 4. The Sun cluster equipment is on a separate electrical circuit allowing for meter readings to be taken specifically related to this system.

The new installation created changes to the HPC topology with the authentication server and NAS relocated to the data centre. A simplified topology is shown in figure 4. As the new installation benefits from a direct connection to the 10Gb backbone, the two network area storage devices were then able to synchronise their data. A further advantage of this topology change was that the authentication server could be used to direct HPC only traffic across the clusters, allowing for larger, potentially more complex jobs requiring multiple processors to be allocated to the new cluster.

As the equipment is now housed in a data centre this has enabled us to have access to power usage data, which resulted in questions about the energy efficiency versus the overall usage of the system. While this is the second HPC installation on campus, it is not possible to compare power usage for both systems. The first and smaller of the two HPC systems is not

located in a data centre, so the actual power usage is not known.




Fig. 4. Simplified data centre topology

### C. Data Transfer Across HPC Network

During the installation of the new cluster and network system configuration, there were no problems encountered with the hardware or software. However, there were issues of the impact on the university network as a result of topology changes, and the need to provide high speed connectivity with as few hops as possible to the network backbone.

There is a dedicated 10 GB fibre link between the HPC and the main campus network. Data transferred across this link is HPC specific. Therefore, application data such as email or web requests are not passed through this link. This also means that the statistics we were able to gather relate directly to HPC traffic with very little extraneous data.

As this is an IP network, TCP will check for correct delivery and resend if a failure occurs. This can be seen from statistical data produced by the switch which shows only two packets having been received with errors and subsequently dropped as shown in Table I.

## IV. NETWORK MONITORING

There is a complexity in monitoring of HPC system which resides within an educational institution. Research activities are not continuous, so there are periods of constant, heavy use and those of very little use. However network monitoring is a useful method to establish a prediction of HPC usage for the next academic year based on maintaining a consistent number of users.

Ganglia network monitoring software [4] is used to monitor the performance of each node of the HPC system. We used Ganglia to monitor the activity across the HPC system over a four month period as shown in Table II. Our primary interest was the amount of data transferred or throughput, between the HPC system and the main campus network. These statistics represent actual data in, out and the total for each month.

TABLE I. ACTUAL DATA TRANSFER DECEMBER TO MARCH 2012

Switch data showing throughput and errors per port				
Total packets received errors	Total packets without errors	Total packets with errors	Broadcast packets received	Packets transmitted without errors
16271883106	2	722158	11752593983	
51931204083	0	28754	12754174868	
339970419	0	339048	227824568	
37714391697	0	244179	138393332334	
662809012	0	263776	409744236	
544547541	0	162119	874139208	
218406251	0	111080	311713912	
9525452	0	28404	38246169	
99872867145	0	1232098	42701373805	
5297798809	0	302201	5598585083	

The specific metrics allowed us to extract bytes in and out. This includes the transfer of data between two NAS servers located in the data centre and the HPC research office, as they synchronise.

TABLE II. ACTUAL DATA TRANSFER AUGUST TO NOVEMBER 2012

Bytes per Month			
Month	Bytes In	Bytes Out	Total Bytes
August	1129179172	1128208207	2257387378
September	4023709673	4038567846	8062277519
October	8412991174	9920932354	18333923528
November	6843640362	6858377800	13702018163

The Ganglia data has helped to identify future use for the next four months. The use of HPC in our institution is not continuous. Instead usage is cyclical following a pattern related to student activities across the year. For example, at the end of term when there are fewer academic staff undertaking research, the HPC activity reduces dramatically.

The predicted data allows us to determine periods of heavy traffic during which the system is closely monitored for any possible problems. The major cause thus far of hardware failure has been overheating due to excessive use. The increasing volume of data and peak usage has also provided important information about the performance of the campus network and enabled us to identify potential bottlenecks as shown in table 3. To address this issue we began our investigation by looking at throughput and power consumption, rather than the complexity of the jobs being submitted.

## V. TREND ANALYSIS AND PREDICTION

To predict the throughput a linear regression was applied to the known values, bytes-in and bytes-out. This produced a further data set which contained the best fit.

The aim of monitoring network throughput was to observe trends in the data to help identify any future issues in the HPC system operation. These trends would then help to identify periods of high or intense activity which might need more



careful monitoring to avoid performance issues and hardware failures. Table 3 shows the actual data trend for four months.

TABLE III. ACTUAL DATA TRANSFER DECEMBER TO MARCH 2012

Bytes per Month			
Month	Bytes In	Bytes Out	Total Bytes
December	1872480335	1872480335	1872480335
January	3595093540	3902042466	5549676017
February	6429131638	6982780716	10318470267
March	6873920589	7673850456	14791963395

When compared to the actual data for the same period of four months, it is evident that there are some deficiencies in using a predictive trend. In our predictions we have used a linear regression model as shown in figure 5. Using this model we were unable to deal with changes such as new users to the HPC and researchers working towards deadlines.

The actual data throughput increased dramatically resulting in a bottleneck which caused significant data loss. If the trend data had been considered earlier, this bottleneck may have been identified sooner and prevented. The end result of this bottleneck was a significant increase in bandwidth between the network area storage servers from 1 Gb to 10 Gb.




Fig.5. Comparison of total data throughput – trend vs. actual over 4 months

The data sets were expanded to include both bytes-in and bytes-out for four months – December to March. The resulting prediction is shown in Figure 6. The predicted data shows a trend that is much closer to the actual throughput. This is a benefit of extending the data set to predict more accurately the throughput and help with identifying possible issues and allowing informed decisions about data transfer and changes to the campus topology.

The campus HPC network infrastructure consisting of geographically distributed clusters provided the essential HPC resources for our expanding research community. However, the power consumption necessary to power and cool these

resources has also increased. In the next section we are considering the power usage related to the network throughput.




Fig. 6. Comparison of data throughput – trend vs. actual over 8 months

## VI. THROUGHPUT VS. POWER

It is widely recognised that HPC systems were not originally developed to be green but to be able to handle high quantities of data [9]. HPC systems are energy hungry. The current trend in measuring HPC and data centre efficiency is to measure either the performance or the energy consumption [10]. Having considered the throughput as a standalone measurement, we realised that in order to assess the efficiency other metrics needed to be included. To investigate this further we utilised Sol to evaluate the relationship between throughput and power consumption. The intention was to use this measurement to predict the power usage in relation to the throughput. This has the potential to develop further into a predictive tool which could help to calculate cost of ownership based on the usage.

Our initial investigation identified that the readings from the electricity meters were not taken at regular intervals, hence the first reading was a cumulative reading for the period August 2012 to June 2013. This meant that the available sample of weekly readings was only available for a short period of four weeks. Whilst this is not a representative data set, it shows some interesting indications of usage.

Our findings demonstrate that there is no real correlation between throughput and power. Instead the relationship between power consumption and HPC usage is deduced from power consumption and the type of jobs running on the HPC, rather than the data throughput. More complex processing requires more CPU, which in turn requires more power. The energy consumption is not dependent on a high throughput. Table 4 clearly shows an increase of throughput between 11<sup>th</sup> and 17<sup>th</sup> June, consuming 2748 kWh.

TABLE IV. POWER CONSUMPTION WITH RELATED THROUGHPUT

Week no	Power Consumption (kWh) with related throughput (bytes)			
	Time Period	Throughput	kWh	kWh Difference
1	August 2012 to June 5 <sup>th</sup> 2013		123735	
2	05/06/13 - 10/06/13	1546596071	12309	1910
3	11/06/13 - 17/06/13	6657466778	51852	2748
4	18/06/13 - 24/06/13	2315395533	17649	2800
5	25/06/13 - 01/07/13	396725019	2962	2755
6	02/07/13 - 08/07/13	164559549	1207	2336

The results show limited to no correlation between throughput and power usage as far actual packet transmissions are concerned. There is a small correlation between the processes caused by the throughput and power usage, however this should be more associated with memory and CPU usage rather than network traffic. There is one argument that the system was performing at a high level of efficiency based on a power consumption of 2748 kWh with a high volume of throughput. There is a further argument which suggests that the efficiency can be measured by looking specifically at the nature of the submitted jobs. The more complex the jobs, the more CPU and memory may be required, the more power is used. Therefore the level of job complexity against power consumption indicates the levels of efficiency. This second argument could prove difficult to assess as every job would have to be logged and examined in detail. This would also require an understanding of the nature of the problem being processed as well as knowledge of the software employed.




Fig. 7. CPU usage showing increase for complex processing

To extend this further we investigated the memory use for the same period. This not only showed an increase in memory use which matched the period of increased throughput, but also for the following period which had considerably less throughput but more intense processing.

Memory and CPU on the other hand by their own nature and architecture increase power consumption when operating at higher loads. This is where kWh can be monitored and observed to increase after the node has received a set of instructions via throughput. At this stage we considered the architecture rather than the level of complexity of jobs.




Fig. 8. Memory usage

#### A. Cumulative Metrics

Each of the metrics considered this far have their own merit in showing the efficiency of their specific function of the system. We identified that CPU, memory, throughput and power consumptions are key components to understanding system efficiency. Consider that a measurement of efficiency can be identified by determining the cumulative cost of these metrics.

$$e = \frac{t + m + c}{k} \quad (1)$$

Where  $e$  is the measured system efficiency,  $t$  is throughput,  $m$  is memory usage,  $c$  is CPU utilisation and  $k$  is kWh. This combination of metrics allowed us to explore a cumulative metric approach. This formula produced an interesting set of results despite the small size of the sample data, as seen in tables V and VI. What this shows is that the higher the system efficiency ( $e$ ) the more efficient the system is as shown in table VI.

TABLE V. CUMULATIVE MEASUREMENT OF EFFICIENCY

Week No	Cumulative measurement data			
	Throughput	Memory	CPU	kWh
2	1546596071	1184969455317	141	12309
3	6657466778	3686159033111	591	51852
4	2315395533	3185239568418	1299	17649
5	396725019	3706923972429	482	2962
6	164559549	3525595150447	420	1207

TABLE VI. CUMULATIVE MEASUREMENT OF EFFICIENCY

Week No	Efficiency as a cumulative measurement		
	Sum of metrics	kWh	Efficiency (e)
2	1186516051530	12309	96392207
3	3692816500480	51852	71218048
4	3187554965250	17649	180611117
5	3707320697930	2962	1251718871
6	3525759710416	1207	2919943809

These results show that too many inconsistencies to draw any reasonable conclusion based on the statistical data. There are high metric totals with low kWh and low metric totals with high kWh. What these results do indicate is a need to consider the nature of the submitted job, the software used and the level of complexity involved in processing.

### B. Other factors impacting energy consumption

The current system provides two network area storage devices which synchronise data across the network. The effect of this is that as throughput increases, the number of read/write increases and the power consumption also increases [8].

This area requires further investigation to understand the effect of data synchronisation on the energy efficiency of the system.

## VII. CONCLUSION

In this paper we have presented our experience in developing HPC system infrastructure at the University of Huddersfield, with a special consideration of the impact on the university network and data centres. All too often HPC systems evolve from being located in a back office to a secure location within the individual departments. That poses a number of issues relating to the data transfer over low bandwidth networks and requires additional provision for power and cooling the HPC equipment. We have analysed the transfer of large data files and highlighted the constraints of a department based infrastructure. As a result of this analysis we have redesigned the campus HPC network to provide high speed connectivity with the university network backbone.

We have demonstrated that the centralised university HPC solution, incorporated into the university data centres, improves the overall utilisation and HPC system performance.

Our experience in integrating new HPC devices into an existing infrastructure will enable us to make better informed network design decisions in the future.

The benefits of network monitoring cannot be overlooked as they play an ever increasing role in understanding the performance and energy efficiency of any HPC system and data centre operation. In this paper we have described our own experiences of developing HPC resources to meet the needs of our researchers. We have shown the importance of data throughput in helping to manage and develop the network topology providing access to local and national HPC resources.

As the energy efficiency of HPC systems and data centres becomes increasingly more important we have shown how the specified metrics can be utilised in understanding system

management, with specific focus initially on data throughput to aid prediction of possible network issues. We then extended the focus to include power consumption and examined the possibility of a correlation between these metrics.

This relationship between power consumption and data throughput led us to look at memory usage, CPU utilisation and to consider the different type of jobs running on the HPC system.

The relationship between data throughput and usage provides a simplistic mechanism to help network and HPC administrators understand the effect that big data has on the network. For HPC this must be considered over at least six months for the data to be of any value. Applying a linear regression to predict trends in throughput does not consider periods of high and low activity.

We have shown that it is not enough to measure the performance of the cluster alone. The increase in large data processing requires the network to respond with sufficient bandwidth to manage predicted throughput. On this basis we have opted not to continue to use a linear regression trend model but to consider the impact of additional metrics alongside power consumption.

With this in mind we have considered the effect of including memory and CPU data and this has led us to believe that further metrics are required in order to successfully measure data centre and HPC system efficiency.

The lessons learned have enabled us to influence the network topology changes to provide a resilient network with sufficient bandwidth to manage our HPC researchers' needs.

## VIII. FUTURE WORK

To extend our work on the energy efficiency of HPC systems within the university data centres we are undertaking further research into data centre infrastructure management (DCIM) and data centre predictive modelling (DCPM). This study will aim to establish the feasibility of designing and building a tool to help data centre managers see through the complexity within the data centre and to select a mode of working which offers more energy efficient compute facilities.

## IX. ACKNOWLEDGEMENT

Our thanks to University of Huddersfield Computing Services department for their support on this project and to all members of the HPC Research Group and University of Huddersfield, Computing and Library Services for their assistance in this installation.

A special thanks to Joanna Radley, Head of Data Centre & Network Services, University of Huddersfield for her help throughout this project.

## X. REFERENCES

- [1] Asvija, B.; Kalaiselvan, K.; Ateet, M.; Sridharan, R.; Krishnamurthy, S.R., "A Performance Based QoS Aware Resource Brokering Framework for the Grid," Information Technology: New Generations (ITNG), 2012 Ninth International Conference on , vol., no., pp.508,515, 16-18 April 2012



- [2] D. Azevedo, A. Hawkins, Z. Ortiz, R. Tiple, S. Van Den Berghe, Compute Efficiency Metric – DCcE, 2010, The Green Grid.
- [3] Chetsa, G.L.T. Lefevre, L. Pierson, J. Stolf, P. and Da Costa, G., Beyond CPU Frequency Scaling for a fine-frained energy Control of HPC systems, " Computer Architecture and High Performance Computing (SBAC-PAD), 2012 IEEE 24th International Symposium on , vol., no., pp.132,138, 24-26 Oct. 2012
- [4] Ganglia Monitoring System, (2001, Jan), [Online] Available <http://ganglia.sourceforge.net/> accessed September 2012
- [5] Hsu, C.-H., et al. (2012). Towards efficient supercomputing: searching for the right efficiency metric. Proceedings of the 3rd ACM/SPEC International Conference on Performance Engineering. Boston, Massachusetts, USA, ACM: 157-162.
- [6] R. Kenway, P. Coveney, S. McGuire, O. Parchment, and M. Parsons, (2012), UK Participation in HPC at the European Level, e-Infrastructure Leadership council, [Online] Available [https://www.gov.uk/government/uploads/system/uploads/attachment\\_data/file/34675/12-1247-uk-participation-in-hpc-at-european-level.pdf](https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/34675/12-1247-uk-participation-in-hpc-at-european-level.pdf)
- [7] Kureshi, I. (2010) Establishing a University Grid for HPC Applications, Published by the University of Huddersfield
- [8] Molka, D. Hackenberg, D. Schöne, R., Minartz, T. Nagel, W., Flexible workload generation for HPC cluster efficiency benchmarking, Computer Science - Research and Development, , Volume 27, Number 4, pp. 235-243, November 2012
- [9] Rone, GE. Cameron, KW. Power-aware high performance computing, in Zomaya, Albert Y. Lee, Young-Choon, Wiley Series on Parallel and Distributed Computing : Energy Efficient Distributed Computing Systems (2012), First Edition, Wiley
- [10] Subramaniam, Balaji, Feng, Wu-chun, GBench: benchmarking methodology for evaluating the energy efficiency of supercomputers, Computer Science - Research and Development, Springer-Verlag, Volume 28, Number 2-3, pp. 221-230, May 2013
- [11] <http://stfc.ac.uk>, Science & Technology Facilities Council, accessed July 2013
- [12] <http://stfc.ac.uk/hartree>, STFC Hartree Centre, accessed July 2013
- [13] Wireshark, [Online] Available <http://www.wireshark.org> accessed September 2012