26th August, 2010

Extreme Computing: Challenges, Constraints and Opportunities

Anne E Trefethen

Preamble

This paper is a handout to support the presentation at the Fujitsu HPC Users Group Meeting on 26th August, 2010.

Abstract

Scientific applications require an ecosystem of computational infrastructure that can allow distributed collaboration and computation, large-scale simulation and analysis and appropriate consideration of data driven science. The ecosystem requires the integration of high-performance computing, in some cases to the exascale, cloud computing, databases, high bandwidth networks and the software and people infrastructure that enables the effective use of the components.

In this presentation we will consider some of the challenges and constraints that drive the development of the computational infrastructure and its components – including computational models, energy and knowhow; and the opportunities that are presented, in terms of new science applications and new algorithmic approaches. There is a particular focus on energy.

Key words

Exascale, extreme computing, energy efficiency, cloud computing

Introduction

Advanced computing is an essential tool in addressing scientific problems of national and scientific interest, including climate change, the virtual human, new materials, next-generation power sources and astrophysics, but as importantly it is equally essential to solve commercial and industrial problems in financial modelling, engineering, and real-time decision systems.

Extreme computing has moved from being a somewhat esoteric interest of a few scientists to a necessity for any computational scientist or developer who uses simulation as a tool. A shift has occurred in recent years as the processor chip designs mimic the architecture of high-performance computers, with multiple processing cores on a single chip, making efficient programming of a single processor computer as complex as it once was to develop software for a high-performance computer.

The interface to extreme computing is also changing with increasing availability of infrastructure and services through Cloud Computing enabling access to both high-performance and high-throughput computing.

The computers and computer systems, however, are still largely designed with little thought to the applications that might use them and often arrive to the users with little software infrastructure in place. The development of the required software infrastructure is becoming increasingly complex and if we are to be effective in the use of these systems there is a need to address this through appropriate abstractions, programming languages and tools to support application developers and users.

A further pressure point for extreme computing is the constraint that energy requirements are having, and will continue to have, on the developments of computer architectures and their use. This becomes increasingly important both in terms of the general computing infrastructure which is becoming a significant component of national use of electricity [1] and for the high-end computer systems, with the predictions of the exascale energy requirements up to 100MW [2] – more than would be required by a small town. Given that a standard dual-core laptop is equivalent to one of the Top500 machines of only 12 years ago we can expect these issues to be pervasive as we go forward.

This presentation provides an overview of the challenges that are faced in extreme computing, together with some of the ongoing activities to provide a roadmap for future developments. Focussing on energy as a driver there is a consideration of some of the research and activities across the ecosystem of computing and in particular we consider that required for the Square Kilometre Array [3] which provides an excellent co-design vehicle[1].

Ultimately our interest is in providing a framework to allow the development of applications and algorithms that are energy-aware and will allow the user to make choices regarding the objective optimized in the software – be it energy performance, price performance or indeed computational performance.

---

[1] A *co-design vehicle* is a concept described in [ref] and is essentially an application that is developed in conjunction with the design of the software and hardware infrastructure together.

Extreme Computing Challenges and Constraints

There have been many studies during the last several years that have addressed the requirements of extreme (or exascale) computing [4, 5]. Here we will provide a general overview with some specifics from a recent study in the UK focussed on the development of a roadmap for high-performance computing [6] and the roadmap of the International Exascale Software Program [7] but the reader is referred to the original texts for the full discussion and details.

The HPC/NA roadmap activity in the UK brought together computational scientists, mathematicians and computer scientists to understand what might be the key issues for the future of extreme computing both in the UK and within an international context. The activity held a number of workshops and had individual communications with scientists and groups in the UK and a review of the literature from other national activities was included to provide the international perspective. It did not address hardware challenges directly, but the whole discussion was of course motivated by the changing face of computing hardware. The IESP roadmap provides a good overview of the hardware expectations [8]. As we go forward, we can expect that power and cooling limits will continue to constrain clock speeds; we will see a shift to multi/many core with attendant hierarchical parallelism an often with the additional complexity of hardware accelerators. This can be seen by the evidence of the Top 500 machines where multi-core machines and those with GPUs and the like are an increasing proportion [9]. Memory bandwidth will continue to be a bottleneck and this is likely to get worse. Chip design and chip to chip communication capabilities are likely to change however with 3D stacking and integration designs, for example [10] and the introduction of high-speed optical communications [11].

The HPC/NA roadmap activity found there were five major themes that dominated the discussions.

Theme 1: Cultural Issues

There are gaps between disciplines that meant that often the team effort that might be required to bring a large-scale application together was not optimal and that appropriate knowledge and expertise might not be captured within such an effort. This is a common issue and many national activities are investing to ensure that community actions are supported and interdisciplinary teams created.

Differences in discipline communities were also identified in that some application domain scientists are used to sharing models and codes, and reusing other software developed by other groups; while for other domains this approach is almost completely alien with codes being entirely developed within a particular group and little use being made of libraries or other third-party software.

Similarly there is a need to ensure that the software activities are not only focussed at a national level but are within an international context. It is the case that several application codes in the UK are dependent up on software that has been developed in programmes in other countries and vice versa. These dependencies are a risk to the sustainability of the software and that risk can only be mitigated by international collaboration.

Theme 2: Applications and Algorithm development

There are many issues that need to be addressed in the area of application and algorithm development. These include ensuring that a component approach to development is taken where appropriate such that algorithms that might underpin multiple application areas are developed as such. Having noted that, many applications involve multiple models at different scales or for different elements of the application and integration of components is not possible without appropriate standards for data models and formats, interoperability of programming models, and lack of knowledge of error propagation through the integrated system. Clearly this is not a problem unique to extreme computing, but is one that as a community we have yet to address fully.

The memory hierarchy of high-performance computers is on course to get even more complex which will drive the needs for hierarchical algorithms to deal with bandwidth across the memory hierarchy together with software strategies to mitigate high memory latencies. This will in turn drive the need for algorithms to be dynamically adaptive, perhaps as components of an active library, and of course scalable much beyond the present status quo. There are few applications that can scale to petaflops levels let alone toward exascale. Such scaling can only be achieved with appropriate portioning and load balancing and effective data management. Data-intensive science [12] is increasingly important and the ability to manage large amounts of data effectively and efficiently is crucial [13]. Moving data takes energy, at every level of the system – this is a point we will return to later.

Theme 3: Software Challenges

Development of software is becoming increasingly complex and without appropriate software engineering, in the context of an exascale system, will be un-maintainable. Application scientists are not usually software engineers but there is a need to ensure that by some means appropriate practice is adopted. Frameworks and tools to support software development will be needed together with compilers and code generation tools that can provide a layer of abstraction for application scientists.

Theme 4: Sustainability

The HPC/NA Roadmapping activity identified a general concern regarding the sustainability of application codes, software libraries and skills (we consider skills in the next section). This issue is integral to that of programming models, interoperability and also the cultural issues above. Scientists are naturally loathed to invest a great deal of effort in the development of software that will last only the lifetime of a particular computer architecture, compiler or other dependent component of the environment. We can only address this issue as a community as we agree in the adoption of standards and open source mechanisms.

Theme 5: Knowledge Base

It is important to ensure that computational scientists have the right set of skills, and this set of skills needs constantly updating. Within the UK activity it was also discovered that there was a lack of awareness by some members of the community of existing libraries/packages. Maintaining a strong knowledge base will require education for graduate students as well as the opportunity for life-long learning. The report on exascale computing for energy and environment [14] notes “The current belief is that the broad market is not likely to be able to adopt multi-core systems at the 1000-processor level without a substantial revolution in software and programming techniques for the hundreds of thousands of programmers who work in industry and do not yet have adequate parallel programming skills."

Other issues:

Since the last of the roadmapping workshops Cloud Computing [15] has matured and is now very much part of the computational environment. It offers a different interface to distributed, and in some cases, high-performance computing and as such is part of the evolving ecosystem for application development. As data-intensive applications dominate it is possible that the business models underpinning Cloud computing might offer benefits to scientists. These issues will likely be resolved as computational scientists become engaged with the various vendors but there is no question that the scale of Cloud computing infrastructure may offer benefits in terms of energy efficiency and from that point of view we will return to them as part of the ecosystem later.

Co-Design Vehicles and the Square Kilometre Array

During the workshops and activities of the Department of Energy in the USA addressing the requirements for exascale computing, the notion of a co-design vehicle has been adopted. At present, other than in a few specific cases [15, 16, 17], a computer system is designed with little thought for the applications that will use it, and the software is generally designed to try to fit the hardware. In the co-design model the two are seen as part of the same system and the design of each is done through collaboration. In the extreme this might result in computer systems that are only useful for a given application such as [15, 16, 17] but the model can be adopted to enable more generic progression with the evolution of the computer architect and algorithms being better aligned and providing better capability and specifically energy performance.

We are interested in the computing challenge provided by the development of the Square Kilometre Array [?], the next generation radio telescope. The telescope is in the design phase with the anticipation that construction of the first phase will begin in 2016 with the full telescope completed and in operation by 2022. The SKA will likely be located in either Australia or South Africa, in a desert so as to have little or no interference, but is a collaborative effort of over 50 groups in 19 countries.

The present design [18,19,20] has a combination of aperture arrays in the core and up to 3000 phased array feeds on dishes and a collecting area of approximately one square kilometre with receptors extending out to a distance of 3000km from the centre of the telescope (figure 1). It will allow a sensitivity of more than 50 times that of existing telescopes, and 10,000 times the survey speed and will provide data to answer fundamental science questions on gravitation and magnetism, galaxy formation and even the question of life on other planets. The design of the SKA is developing through design studies based on the science requirements, Pathfinder telescopes that provide experience of design options, and technology capability considerations.

Figure 1 Possible configuration of SKA receptors and artist's impression of SKA core(from [21])

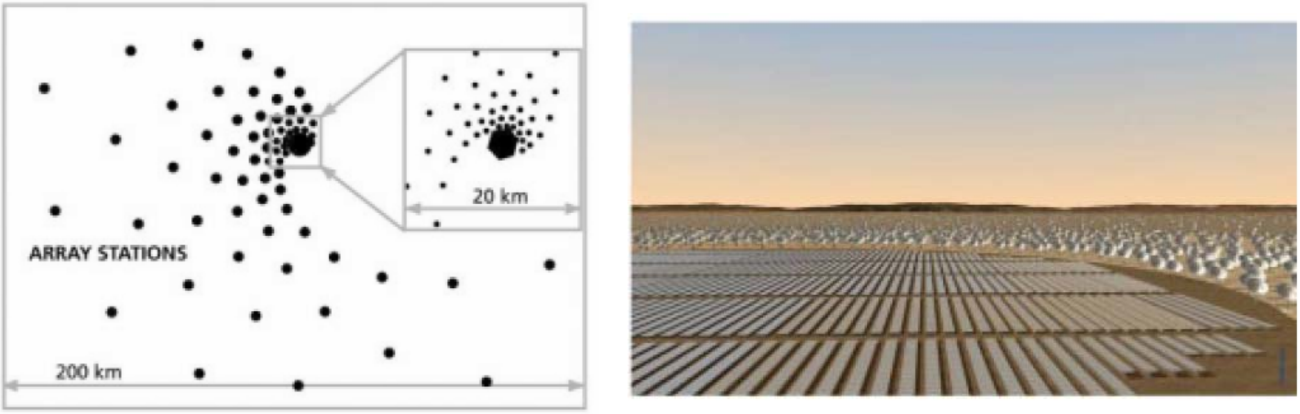

The SKA provides a fabulous information technology challenge with a typical data rate from each dish antenna on the order of 100Gbs^{-1} aggregating to over 100Tbs^{-1} [21] and need for Exaflop computation [ref] for post-processing. The infrastructure required to support the various science cases will need to range from real-time capability to transport and analyse the data at these high-data rates and the capacity to store and "publish" the data for later analysis and interpretation by the global astrophysics community. The computational systems will likely range from specifically designed FPGA-like units to exascale computing and Cloud-like data centres. The communications infrastructure will range from intra-chip and inter-chip with optical fibre to the correlator and on to a high-performance computer, to trans-oceanographic with the latter having data rates of at least 100Gbs^{-1} over the general network providers. The SKA will succeed or not depending on both the physical implementation of the telescope design and the software infrastructure that will enable it. The software infrastructure required to realise this information technology challenge is itself has been identified as > 2000 person year task [22] but even this may not take full account of the complexity of the task.

Figure 2: SKA computational requirements in context of Top500 [23]

The fundamental algorithms in the SKA data pipeline include beamforming, gridding, convolution and filtering algorithms. Cornwell et al [23] estimate the computational requirements in the context of the Top500 and show them to be beyond the scale of projected performance in the timescale required (figure 2).

Alongside the challenges of the computational infrastructure are the related challenges of powering the infrastructure. This includes the power required for the core antennas (~30MW) and the remote stations ~0.5MW) and the various computational components including high-performance computing (~40MW) data transmission (~?MW) and Cloud (or equivalent) provision (~?MW). Energy provision is a major design factor in the delivery of the telescope – plans to mitigate the energy constraints include renewable energy sources (sun) providing the station power and it is easy to see that the SKA could possibly be the largest Green IT project ever to be considered.

Counting Joules

As indicated above the ecosystem of computational resources required to enable the SKA provides a plethora of challenges and opportunities where energy efficiency is concerned.

Moving data takes energy whether the data is moving from L2 cache on chip or within a transatlantic Data Cloud. Effective management of that data communication is a major component of any optimal energy model. There has been a great deal of research on wireless network communications and sensor networks, where devices are most often low-energy devices with battery constraints. Indeed a lot of research has been done in general on low-energy devices including computational algorithms from which we might learn. The existing communications network across the globe relies on hundreds of thousands of switches and routers and these, unlike wireless or mobile infrastructure, do not power down when idle. The network is a complex system of different technologies and it is difficult to predict the energy consumption under different circumstances so sending data 10 times as fast might use interfaces that use 100 times the amount energy or in some cases (e.g. newer optical ones) 1000 times less [24]. A good overview of network energy costs for realistic configurations can be found in [25]. There are a number of research efforts considering the issues around Green network communications including INTERNET [26] and others can be seen at recent conferences [27,28]. Of course as energy-aware systems are developed there are still issues of what they optimise – usage or cost? Qureshi et al [29] illustrate effective ways in which energy-aware data centres can optimise cost by moving computation to nearby states where electricity costs are less.

Optimizing energy usage in large-scale data centres and Clouds is almost a science in itself. McKinsey and the UpTime Insitute [1] indicate that the energy used by data centres in the US is becoming a significant percentage and is likely to overtake airlines in terms of carbon emissions. The report states that the average data centre uses as much power as 25,000 households, but that estimate is probably somewhat out of date as in the last few years Cloud provisioners have built very –large scale data centres across the US [30]. On a somewhat smaller scale at the Oxford Supercomputing Centre (OSC) we have developed software that intelligently powers down components of the systems at times of under utilisation. The indications are that this will provide significant savings.

The McKinsey report identifies a number of issues around the effectiveness of data centres including siloed organisations and limited transparency that match very well with the findings of the HPC/NA for that community. McKinsey make the recommendation that metrics be defined that are not only measuring the facility but are linked to the applications using it and the processes integral to it. A consortium called the Green Grid is now in place with the aim to develop standards and best practice for data centres. The recommendation regarding metrics is one that we, as a community, should also take on board as we develop exascale technology.

At the computing system level there are many approaches to energy efficiency. Again there is a large knowledge base in the low-energy systems community that we should consult. Indeed an approach that might be considered for some application areas builds directly on the low-energy embedded microprocessor technology to build petascale systems. This approach is under development for the Green Flash system [31, 32] at Berkeley National Laboratory where researchers are collaborating directly with Tensilica, Inc. to explore the use of Tensilica's Xtensa processor cores in building a computer to model clouds (that is clouds in the weather system not as a computer system!). They believe that a specially designed core could get 10–100 times better performance per watt. The Green Flash design will have on the order of 20 million processors consuming less than 4 MW whereas the equivalent in conventional microprocessors would require around 200MW.

With appropriate systems support it is possible to provide hooks into the operating system to allow spindown policies to be enacted, adaptive placement of memory blocks, agile use of component devices and of course energy-aware routing [33].To make use of this level of support it would be helpful to have tools that can provide a mapping from the higher-level action to the underlying activity.

The power cost for a device is proportional to the frequency cubed. This means that a multi-core device that is running at a lower clock rate is bound to provide the potential for higher performance at a lower energy cost. Interesting issues arise when there are multiple such devices to hand and the algorithm designer can chose to optimise the energy usage for a given computational performance. At this point in time the SKA community are considering the use of GPUs, FPGAs and multi-core chips [34, 35]. Nieuwpoort et al [34] show some interesting results in this area where they have used a variety of chipsets to underpin the computation to correlate signals. The results are particularly interesting in respect to the achieved efficiency vs Gflops/Watt where although a given device might underperform in terms of percentage efficiency it may still perform better in Gflops/Watt. The Green500 provide good reference indicators for feasible platforms at this time [36].

Within the Oskar project [37] we tackled the Digital beamforming for the aperture array components of the SKA that pose considerable computational challenges [38]. The proposed algorithm provides a hierarchical algorithm for beamforming using a simplified and flexible computational approach of direct matrix-vector multiplies rather than FFTs that provides a reduced data rate and a computational cost of forming beams of 1TFlop as opposed to 20TFlops for the FFT approach.

Within PrepSKA [39] Savlini et al [40] have developed a Pipeline for Extensible Lightweight Imaging and Calibration (PELICAN). This framework for parallel quasi-real time data processing offers two deployment options either with the server supplying multiple pipelines or the pipeline connecting to the data stream directly. PELICAN allows reuse of modular components and will be deployed on LOFAR [41] interferometer stations to allow all sky calibration and imaging, and pre-processing for pulsar searching. The modular framework allows appropriate computational components to be implemented on appropriate devices, in this case GPGPUs.

Tools, Benchmarks, metrics

While benchmarks for datacentres and systems are now relatively mature [42, 43] there is still work to be done at the numerical algorithms level. On the exascale roadmap [8] it is suggested that standards to support energy-aware algorithms should be agreed in 2014/2015 with energy-aware libraries available in 2016. We have our work cut out!

Although for the SKA specific algorithms of interest we can model and implement across different platforms there is a lack of agreed metrics, tools and benchmarks to provide the community with clear evidence to support design features. This is true across the board for computational algorithms although there are some efforts in this direction.

There are a number of tools for a variety of platforms that will allow measurements of device and sub-device power usage. Kirk Cameron and his group have created a profiling tool that is an open source software system together with hardware power measurement devices, PowerPack [44]. Using the profiler they have analysed the HPC Challenge Benchmarks [45] and been able to provide very clear analysis of the behaviour of the algorithms in terms of energy used.

It may be time to consider again what benchmarks should stress. As we consider adaptive algorithms that are able to autotune to heterogeneous sets of processors we may need to consider a different set of characteristics and metrics that are captured by existing suites. There may be instances where replacing data movement by computation provides a saving in energy while maintaining computational speed.

Conclusions

Energy constraints provide one of the major challenges for extreme computing in the future. There are opportunities at every level of the computational infrastructure to address the challenges through both efficient physical hardware but also through more sophisticated use of the physical infrastructure. An intelligent software infrastructure could reduce the energy use on chip, within a processor, between processors, between computers, across Cloud platforms and indeed over international boundaries.

Ultimately our own interest lies in the development of energy-aware algorithms for extreme computing. We believe that to achieve such algorithms and to see the benefits there will need to be significant investment in the support of tools and low-level mechanism to allow the profiling of energy use across platforms. Equally important are standard metrics and computational benchmarks that will allow agreement on measures of "success". Just as the linpack benchmark has come to be seen as the measure of the performance of a computer system we need the equivalent application/algorithm benchmarks that capture the energy characteristics of any given system.

The future holds continuing complexity in computing systems from the combinations of computing within the whole ecosystem to the heterogeneity of chips on our laptops. The most effective algorithms at any given point in time for any specific application are going to depend upon the choices that can be made given the hardware configurations. In general, at present we cannot even articulate those choices let alone provision the most appropriate algorithm. This challenge of articulation and optimal algorithm design can only be addressed through a co-design approach with hardware, software and application scientists working together.

While we have focussed here on the energy aspects of the problem, and specifically as driven by SKA, there are any one of several other key issues that could and should be addressed including the complexity of the software development for such an infrastructure, the usability of the software systems, the data management and related semantic issues for such colossal data systems or the provision of imaging and data analysis in the Cloud. Each of these is a talk of its own – for next time.

Acknowledgements

My thanks to colleagues on the prepSKA project Steve Rawlings, Aris Karastergiou, Stef Salvini, Ben Mort, Chris Williams, Fred Dulwich and Andy Faulkner. The HPC/NA project was in collaboration with Nick Higham, Iain Duff, Peter Coveney, Mark Hylton and Stef Salvini. I am grateful to Jeyan Thiyagalingam, Simon McIntosh-Smith, Jon Crowcroft and Jaafar Elmirghani for their input and suggestions.

References

1. Revolutionizing Data Center Efficiency, McKinsey Company and the Uptime Institute, *www.**mckinsey**.com/.../Revolutionizing_Data_Center_Efficiency.pdf*
2. www.exascale.org
3. The SKA project: www.skatelescope.org/
4. The DARPA Exascale Study Report: http://users.ece.gatech.edu/mrichard/ExascaleComputingStudyReports/exascale_final_report_100208.pdf
5. The Department of Energy workshop on extreme computing: http://extremecomputing.labworks.org/index.stm
6. The HPC/NA Roadmap: www.oerc.ox.ac.uk/research/hpc-na/roadmap
7. The International Exascale Software Project: www.exascale.org
8. The IESP roadmap: http://www.exascale.org/mediawiki/images/4/42/IESP-roadmap-1.0.pdf
9. The Top500 machines: www.top500.org
10. Intel http://www.eetimes.com/electronics-news/4204959/Intel-Silicon-Optics
11. 3D DRAM Design and Application to 3D Multicore Systems, Hongbin Sun, Jibang Liu, Rakesh S. Anigundi, Nanning Zheng, Jian-Qiang Lu, Kenneth Rose, Tong Zhang, IEEE Design and Test of Computers**,** vol. 26, no. 5, pp. 36-47, September/October, 2009.
12. The Fourth Paradigm: Data Intensive Research,http://research.microsoft.com/en-us/collaboration/fourthparadigm/
13. The Data Deluge: An e-Science Perspective, A.J.G. Hey and A.E. Trefethen, in *Grid Computing - Making the Global Infrastructure a Reality*, Berman, Fox, Hey, eds., pp 809-824, Wiley and Sons, 2003.
14. DOE Report on Modeling and **Simulation** at the Exascale for Energy and the Environment, June 2007
15. Above the Clouds: A Berkeley View of Cloud Computing, Michael Armbrust, Armando Fox, Rean Griffith, Anthony D. Joseph, Randy Katz, Andy Konwinski, Gunho Lee, David Patterson,

Ariel Rabkin, Ion Stoica, and Matei Zaharia, UC Berkeley Reliable Adaptive Distributed Systems Laboratory . http://radlab.cs.berkeley.edu/ February 10, 2009

16. The QCDOC project: *ukqcd.epcc.ed.ac.uk/community/qcdoc/*
17. The MD-Grape3 machine: http://www.peta.co.jp/index-en.html
18. The D.E.Shaw Machine*:* www.deshawresearch.com/
19. The Square Kilometre Array, http://www.skatelescope.org/
20. The Square Kilometre Array, Peter E. Dewdney, Peter J. Hall, Richard T. Schilizzi, and T. Joseph L. W. Lazio, Proceedings of the IEEE | Vol. 97,No. 8, August 2009
21. The Square Kilometer Array (SKA) Radio Telescope: Progress and Technical Directions, P.J. Hall, Richard T. Schilizzi, and T. Joseph L. W. Lazio, U.R.S.I, The Radio Telescope Bulletin, No 326, September 2008
22. SKA Memo 100: Preliminary Specifications for the Square Kilometre Array, R. T. Schilizzi, P. Alexander, J. M. Cordes, P. E.Dewdney, R. D. Ekers, A. J. Faulkner, B. M. Gaensler, P. J. Hall, J. L. Jonas, K. I. Kellermann, *http://www.skatelescope.org/PDF/memos/100_Memo_Schilizzi.pdf*
23. Scaling Mount Exaflop: from the pathfinders to the Square Kilometre Array, T.J. Cornwell, *Ger van Diepen, www.atnf.csiro.au/people/tim.cornwell/publications/MountExaflop.pdf*
24. Green Optical Communications—Part II: Energy Limitations in Networks, Rodney S. Tucker*,* IEEE JOURNAL OF SELECTED TOPICS IN QUANTUM ELECTRONICS
25. Private communication with Jon Crowcroft
26. First ACM SIGCOMM Workshop on Green Networking, August 2010, *http://conferences.sigcomm.org/sigcomm/2010/gncfp.php*
27. INTERNET: Intelligent Energy Aware Networks, *http://www.internet-project.org.uk/*
28. E-energy: First international conference on energy-efficient computing and networking, April 2010, *http://www.e-energy.uni-passau.de/*
29. Cutting the Electric Bill for Internet-Scale Systems, A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag and B. Maggs. SIGCOMM 2009.
30. Data Centres in the News: *http://www.datacenterknowledge.com/archives/2010/02/05/virginia-nc-battling-for-microsoft-data-center/*
31. The Green Flash computer project, *http://www.lbl.gov/cs/html/greenflash.html*
32. Towards Ultra-High Resolution Models of Climate and Weather , Michael Wehner, Leonid Oliker, and John Shalf, *International Journal of High Performance Computing Applications May 2008 22: 149-165, doi:10.1177/1094342007085023*
33. Every Joule is Precious: The Case for Revisiting Operating System Design for Energy Efficiency, Amin Vahdat, Alvin Lebeck, and Carla Schlatter Ellis, Proceedings of the 9th workshop on ACM SIGOPS European workshop: beyond the PC: new challenges for the operating system Pages: 31 – 36, 2000 ISBN:1-23456-789-0
34. Using many-core hardware to correlate radio astronomy signals, Rob V. van Nieuwpoort, John W. Romein, Proceedings of the 23rd international conference on Supercomputing, pages: 440-449, 2009, ISBN:978-1-60558-498-01
35. Evaluating Multi-Core Platforms for HPC Data-Intensive Kernels, Alexander S. van Amesfoort, Ana L. Varbanescu, Henk J. Sips, Delft University of Technology, The Netherlands
36. The Top Green 500 Computers, *www.green500.org*
37. The OSKAR project: http://www.oerc.ox.ac.uk/research/oskar

38. OSKAR: Simulating Digital Beamforming for the SKA Aperture Array, Fred Dulwich, Benjamin J. Mort, Stefano Salvini, Kristian Zarb Adami, and Mike E. Jones, Widefield Science and Technology for the SKA, SKADS Conference 2009
39. Preparation for SKA: *www-astro.physics.ox.ac.uk/~sr/prepska.html*
40. The Pelican framework: *https://wiki.oerc.ox.ac.uk/svn/pelican/slides/pelican-slides2.ppt*
41. The Lofar Pathfinder: *www.lofar.org/*
42. The SPEC power benchmark, *http://www.spec.org/benchmarks.html#power*
43. JouleSort: A Balanced Energy-Efficiency Benchmark, Suzanne Rivoire, Mehul A. Shah, Parthasarathy, Ranganathan, Christos Kozyrakis, SIGMOD'07, June 11–14, 2007, Beijing, China
44. PowerPack: Energy Profiling and Analysis of High-Performance Systems and Applications , Rong Ge, Xizhou Feng, Shuaiwen Song, Hung-Ching Chang, Dong Li, Kirk W. Cameron, IEEE Transactions on Parallel and Distributed Systems, 23 Apr. 2009. IEEE computer Society Digital Library. IEEE Computer Society, <http://doi.ieeecomputersociety.org/10.1109/TPDS.2009.76>
45. Energy Profiling and Analysis of the HPC Challenge Benchmarks Source, Shuaiwen Song, Rong Ge, Xizhou Feng, Kirk W. Cameron International Journal of High Performance Computing Applications Volume 23 , Issue 3 (August 2009), Pages: 265-276, 2009