SPECIAL ISSUE PAPER

The 4 Pillar Framework for energy efficient HPC data centers

Torsten Wilde · Axel Auweter · Hayk Shoukourian

Abstract Improving energy efficiency has become a major research area not just for commercial data centers but also for high performance computing (HPC) data centers. While many approaches for reducing the energy consumption in data centers and HPC sites have been proposed and implemented, as of today, many research teams focused on improving the energy efficiency of data centers are working independently from others. The main reason being that there is no underlying framework that would allow them to relate their work to achievements made elsewhere. Also, without some frame of correlation, the produced results are either not easily applicable beyond their origin or it is not clear if, when, where, and for whom else they are actually useful. This paper introduces the "4 Pillar Framework for Energy Efficient HPC Data Centers" which can be used by HPC center managers to wholistically evaluate their site, find specific focus areas, classify current research activities, and identify areas for further improvement and research. The 4 pillars are: 1. Building Infrastructure; 2. HPC Hardware; 3. HPC System Software; and 4. HPC Applications. While most HPC centers already implement optimizations within each of the pillars, optimization efforts crossing the pillar boundaries are still rare. The 4 Pillar Framework, however, specifically encourages such cross-pillar optimization efforts. Besides introducing the framework itself, this paper shows its applicability by mapping current research activities in the field of energy efficient HPC conducted at Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities to the framework as reference.

Keywords Energy efficiency · HPC · Data centers · Framework · 4 Pillars

1 Introduction

Rising energy costs, and the increase in data center power consumption driven by an ever increasing demand for data services, are becoming a dominating factor for the Total Cost of Ownership (TCO) over the lifetime of a computing system. Additionally, current semiconductor technology will be hitting a point where downsizing and, thus, inherent reduction of power, will no longer be possible mainly due to the increase in leakage currents [23].

Many high profile studies show the increase of data center power demands as well as the power challenges that high performance Exascale computing will provide. For example, a study sponsored by the German Federal Ministry of Economics and Technology (Bundesministeriums für Wirtschaft und Technologie, BMWi) in 2009 [20] showed that data centers in Germany used 9 TWh of electrical energy per year in 2007 with the prediction of 9.5 TWh in 2010 (corresponding to roughly 1.3 Billion Euro at that time), 10.3 TWh by 2015 and 12.3 TWh by 2020. The same trend is visible for the world wide data center power consumption from the years 2000 to 2005 [25] and from 2005 to 2010 [26].

With one eye towards future developments the High Performance Computing (HPC) community realized that energy efficiency improvements are needed for sustainable

T. Wilde (✉) · A. Auweter · H. Shoukourian
High Performamece Systems, Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities (BAdW-LRZ), Bolzmannstraße 1, 85748 Garching bei München, Germany
e-mail: torsten.wilde@lrz.de

A. Auweter
e-mail: axel.auweter@lrz.de

H. Shoukourian
e-mail: hayk.shoukourian@lrz.de

multi-peta and Exascale computing. The "Exascale Computing Study: Technology Challenges in Achieving Exascale Systems" [24] sets a definite power challenge of 20 MW which became a main driver for HPC energy efficient research. Therefore, reducing the energy consumption and overall environmental impact of data centers has become an important research area.

The last couple of years have seen a rapid increase in projects and publications covering a multitude of topics in this field. Unfortunately without some underlying basic structure it is difficult for data center operators to identify center specific optimization approaches, to find relevant research areas and publications, and to identify and apply possible solutions. Current focus areas for data center energy efficiency improvements are: energy efficient application design (which in most cases are directly related to performance improvement efforts like AutoTune [3]); application co-design (developing an HPC system that supports the application behavior in terms of either communication, memory pattern, or I/O requirements, which is the focus of Exascale projects such as DEEP [10] and MontBlanc [30]); batch scheduler controlled Dynamic Voltage and Frequency Scaling (DVFS, like in [15, 19]); and scheduling for a federation of data centers that has been investigated, for example, by the Fit4Green [12] and All4Green [1] projects. Data center instrumentation is another important research area which is concerned with defining how and where to measure certain key performance indicators (KPI's) related to energy efficient HPC. However, current systems often miss proper instrumentation and, therefore, do not allow for the easy collection of required data. The Energy Efficient HPC working group [17] is one of the major driving forces trying to help the HPC community to define the needed HPC system and data center instrumentation and provide guidelines on how to measure the KPI's defined by other standardization bodies such as ASHRAE [2] and The Green Grid [16].

The 4 Pillar Framework for energy efficient HPC data centers was specifically developed to address the lack of a foundation for energy efficiency efforts in high performance computing and to help the HPC data centers to identify areas of improvement and develop applicable solutions. It is completely data center independent by providing a generic way to look at the energy efficiency improvement domain. Other related works share some ideas with the 4 Pillar Framework but leave out significant parts. For example, "Energy Efficiency in Data Centers: A new Policy Frontier" [28] is referring to parts of Pillar 1 and 2 only and, thus, it is solely focusing at the data center energy efficiency from an operational point of view. Another example, the "DPPE: Holistic Framework for Data Center Energy Efficiency" [8], references pillar 1, 2, and 3 to define the data center energy flow. This flow is then used to find areas of improvement and the corresponding division (operating unit) in the data center. It

considers each 'pillar' as a separate improvement area with their own KPI (GEC—Green Energy Coefficient, PUE—Power Usage Effectiveness, ITEE—IT Equipment Energy Efficiency, ITEU—IT Equipment Utilization) that can be measured. This approach has also a strong operational focus.

Both examples can be seen as one specific implementation of the 4 Pillar Framework. They show that the framework can be the foundation for creating energy efficiency models for specific data centers and for modeling data center energy flow chains. This is possible because the 4 Pillar Framework acknowledges that different data centers have different goals and requirements, that applications play an important part for the energy efficiency of HPC data centers, and that the addition of cross pillar interactions allow for more fine-tuned energy efficiency related decisions.

This paper will present the basis for the 4 Pillar Framework, cover the definition and goal of each pillar, and provide a representative model using activities performed at Leibniz Supercomputing Center of the Bavarian Academy of Sciences and Humanities (BAdW-LRZ).
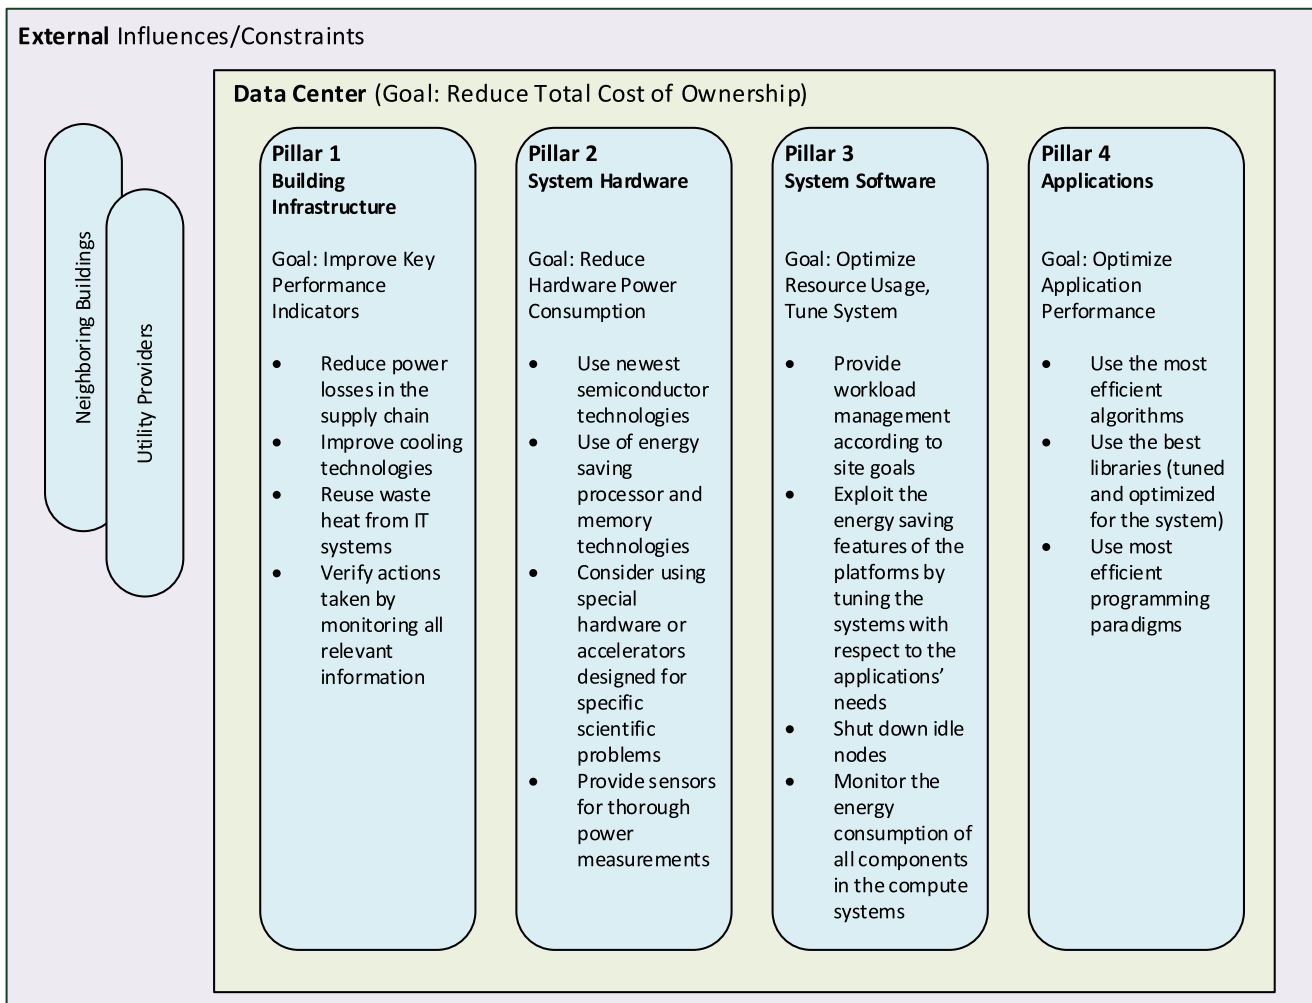
2 The 4 Pillar Framework

The 4 Pillar Framework was developed at BAdW-LRZ mainly to solve two problems. The first was to understand and categorize all internal and external efforts in the area of improving the data center energy efficiency. The second was to provide a foundation for planning future work and for presenting current efforts to outside stakeholders. By addressing these two areas it is possible to work on the BAdW-LRZ data center goal of reducing the TCO for the lifetime of its HPC systems.

The 4 pillars originated from the need to answer the question: What HPC data center aspects play an important part for the improvement of energy efficiency? To answer this question, BAdW-LRZ looked at current research efforts, industry standardization bodies, their own goals and efforts, and at the building blocks of its data center. Using this methodology, four clusters of activities could be identified that make up the 4 Pillar Framework. The identified pillars are:

1. Building Infrastructure
2. System Hardware
3. System Software
4. Applications.

As can be seen in Fig. 1, the four pillars present individual goals and propose different activities to achieve these goals. Yet, one topic that spans all pillars is monitoring. For now, monitoring demands can be summarized for each pillar with 'the more the better'. Without extensive monitoring, no

Fig. 1 The 4 Pillar Framework for energy efficient HPC data centers

system analysis can be done and without analysis, no optimization and, ultimately, no energy efficiency improvements are possible.

2.1 Pillar 1: Building Infrastructure

Pillar 1 "Building Infrastructure" represents the complete non-IT infrastructure needed to operate a data center. The goal of this pillar is the improvement of data center energy efficiency key performance indicators such as Power Usage Effectiveness (PUE), Water Usage Effectiveness (WUE), Carbon Usage Effectiveness (CUE), Energy Reuse Effectiveness (ERE), etc. These have been addressed in research work as well as by many standardization and industrial bodies. Some of the most noteworthy standardization organizations related to Pillar 1 are ASHRAE [2] and "The Green Grid" [16], which have developed best practices guides and evaluation matrices for efficient data center operation. Important pillar aspects are the reduction of losses in the energy

supply chain, the use of more efficient cooling technologies and cooling practices, and the re-use of generated waste heat (which is possible, for example, when using hot water cooling [6]). It is expected that in particular the last item, energy reuse, will become an important part of any sustainable Exascale computing strategy.

2.2 Pillar 2: System Hardware

Pillar 2 "System Hardware" represents all IT systems, networks, and storage systems in the data center. The goal of this pillar is the reduction of hardware power consumption. It is a pillar where data center managers have only minor direct influence. The main instrument for change is the contract phase for new system acquisitions. Yet, even then, one is bound by the available technologies. Fortunately this is also an area where the data center is not required to actively be involved because vendors see efficiency improvements as a big competitive edge. For example, Intel's Sandy

Bridge technology introduces many new energy saving features [29], IBM is actively researching the re-use possibilities offered by hot water cooled supercomputers [35], and Intel is investigating dynamic voltage and frequency scaling for memory [9]. Newer generations of data center hardware usually come with higher efficiency and better monitoring capabilities and control functions. While hardware is constrained by vendor product availability, there are ways to stimulate innovation. For example, by combining system costs with operational (energy) costs into one acquisition budget for a contract, the vendor will be actively looking at how to develop products that are more energy efficient.

2.3 Pillar 3: System Software

Pillar 3 "System Software" represents all system level software that is required to use the system hardware. The goal for this pillar is to make the best use of the available resources and to fine tune the software tools for the system(s). This can be done by using a workload management system, according to the data center policies and goals, by taking advantage of energy saving features provided by the underlying hardware with respect to the application needs [4, 22], and by reducing the number of idle resources whenever possible [31]. Also, the capabilities of the system software are determined by the underlying hardware but, at the same time, some hardware functions are only possible with higher level software support. For example, without the support of P-states in the operating system, they cannot be used even though the CPU hardware supports them.

2.4 Pillar 4: Applications

Pillar 4 "Applications" represents all user applications. The goal for this pillar is the optimization of an application's performance related to specific hardware in terms of raw performance and scalability [5]. This is usually done by selecting the most efficient algorithms, the best libraries for a particular architecture, and the best suited programming paradigms and languages. Any improvement of an application's performance will, in most cases, improve its overall energy consumption. This behavior can be seen with the PRACE 1st Implementation Phase prototypes [21]. Nevertheless, the capability to correlate power consumption with application behavior is an important step toward an understanding of how implementation choices effect the HPC system power consumption [27, 34].

3 Pillar interactions and existing research gaps

Pillars 1 and 2 define the foundation for the data center which, by its nature, is not easily changed. Pillars 3 and 4

build upon the foundation and constitute an area in which the most energy efficient technology gains can be easily developed.

The first step of using the 4 Pillar model for efficiency improvements in the data center is to look at each pillar as being independent from the others. In fact, this is how most data centers are optimized today.

In Pillar 1 the building infrastructure is designed for a specific cooling capacity and in most cases it is optimized for the maximum expected load.

In Pillar 2 newer hardware technology will usually provide the 'current best' in energy efficiency for the current data center restrictions and requirements, for example, the cooling technologies selected for new systems will depend on the data center capabilities in Pillar 1. But there is also an opportunity for further investigations, for example, in the area of network power consumption. Currently the prevailing HPC networks provide no advanced power saving modes. Also the influence of the application communication pattern on the system power consumption is not well understood and requires further investigation [18].

In Pillar 3 the operating system (OS) is an important part of the energy saving possibilities. As most HPC systems use a Linux kernel, which has an active community providing updates, the HPC managers do not have to focus on developing new energy saving features (i.e. the OS updates become a pipeline for these features). The batch scheduling system, which controls the system resource allocation and workload management, has the highest customization potential in Pillar 3. For example, many available scheduling systems can be optimized via pre-defined policies. Additionally, one could implement software based support for specific energy saving policies in the scheduling system itself or as a higher level tool on top of the scheduler. This higher level tool could, for example, validate the workloads energy requirements in order to assure that certain energy-driven policies are not violated before the workload is actually scheduled.

In Pillar 4 the user applications are performance optimized for a specific system; this could include increasing the scalability of the code or just the use of specific compiler flags for the new architecture. In some cases, hand-tuned libraries and compiler intrinsics are used to further optimize the applications for the given platform.

Optimization which focuses on a single Pillar is a valid approach. Effort can then be focused on the Pillars that are the easiest to optimize for individual data centers. However, current results and experiences at BAdW-LRZ show that this approach is hitting an upper boundary in terms of achievable energy efficiency. In most cases further improvements are only possible if optimization tools can use insight gained from other pillars which will result in more prudent decisions.
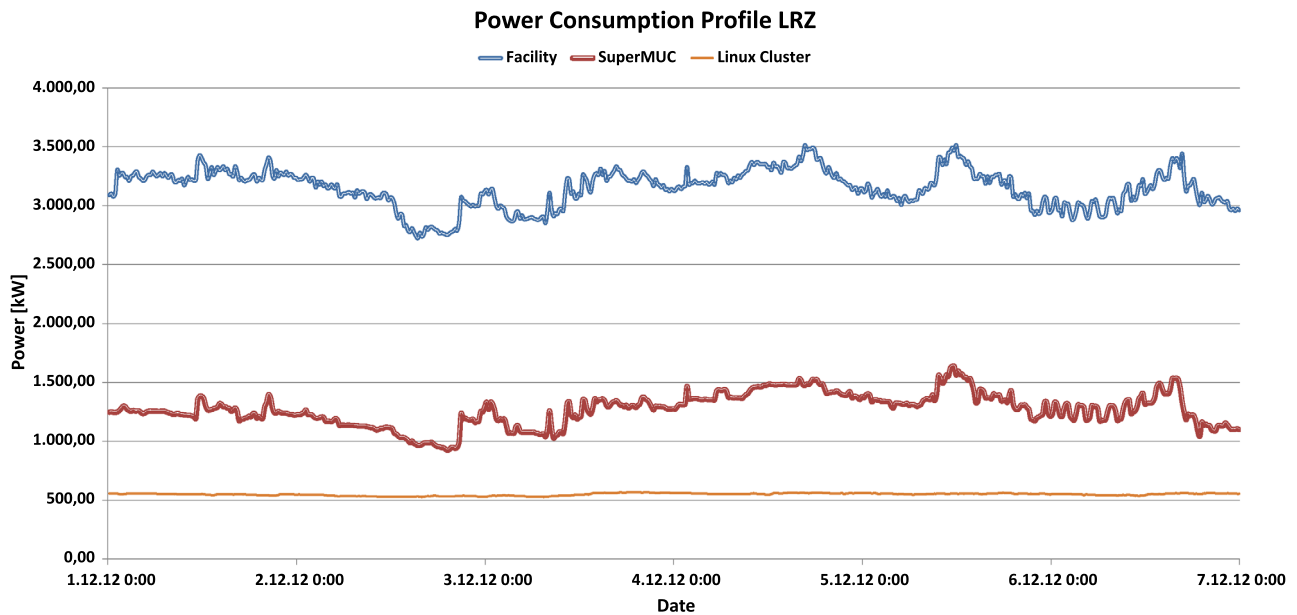
Fig. 2 Power consumption profiles of BAdW-LRZ

A good example showing the necessity of cross-pillar interactions can be given when looking at SuperMUC, the latest supercomputer deployed at BAdW-LRZ. SuperMUC consists of over 150,000 CPU cores in 9400 compute nodes, more than 320 TB RAM, and serves as a Tier-0 supercomputer in PRACE [32]. The power profile of SuperMUC is radically different from previous generations of super computers. Its power consumption varies greatly over time and, consequently, SuperMUC puts a strongly variable heat load on the cooling infrastructure. Figure 2 shows this behavior. The top curve shows the complete data center power draw and, as can be seen, the SuperMUC power draw (middle curve) is dictating the overall power profile. In comparison, one of the older cluster systems (bottom line) shows a nearly constant power draw. This change in behavior is a result of the SuperMUC hardware energy efficiency. Unfortunately, the infrastructure is still optimized for a relative constant power draw. This is one area requiring research.

Another Pillar-spanning concept is heat re-use. Assuming that, for example, in order to heat an office building a water temperature of 40 °C is required, then the outlet temperature of the cooling system (Pillar 1) needs to be kept at least 40 °C. However, the generated heat of the system is in close correlation with the workload profiles. If the work load power draw goes down, a means of retaining the desired outlet temperature could be the deactivation of hardware energy saving features (Pillar 2). On the other hand, if there is less demand for the heat because of warmer weather, the power saving features of the system should be re-enabled.

A last example for a Pillar-spanning scenario is power capping. Here the goal is to control the workload and system resource usage depending on some restrictive data center policies (an example can be found in [11]). These restrictions could be infrastructure limitations (e.g. during a cooling tower maintenance) or business related constraints (e.g. a predefined monthly allowance for power costs). The resource scheduling system (Pillar 3) would require information from Pillar 1 (e.g. available cooling capacity), Pillar 2 (e.g. power consumption of the HPC system hardware with different energy saving techniques) and Pillar 4 (e.g. estimated consumption of applications). Only if this information is available can the scheduler provide support for a specific power capping goal.

4 Application of the 4 Pillar Framework at BAdW-LRZ

Figure 3 shows the current high level activities at BAdW-LRZ. BAdW-LRZ is concerned with heat reuse solutions, collecting data for power monitoring of applications, tuning HPC systems for energy efficiency, and is actively involved in the development of application tuning tools

Figure 4 shows the current activities for each pillar at BAdW-LRZ and the groups responsible for work in each pillar. The CS (Central Services) group is in charge of the BAdW-LRZ data center infrastructure, IBM is responsible for everything related to the SuperMUC hardware, the HPC (High Performance Computing) group is responsible for HPC system administration and low level tools, and the APP (Applications) group is responsible for scaling and performance improvements of specific HPC applications. This view provides an easy way to identify the data center groups that need to coordinate with each other concerning a specific cross pillar activity.
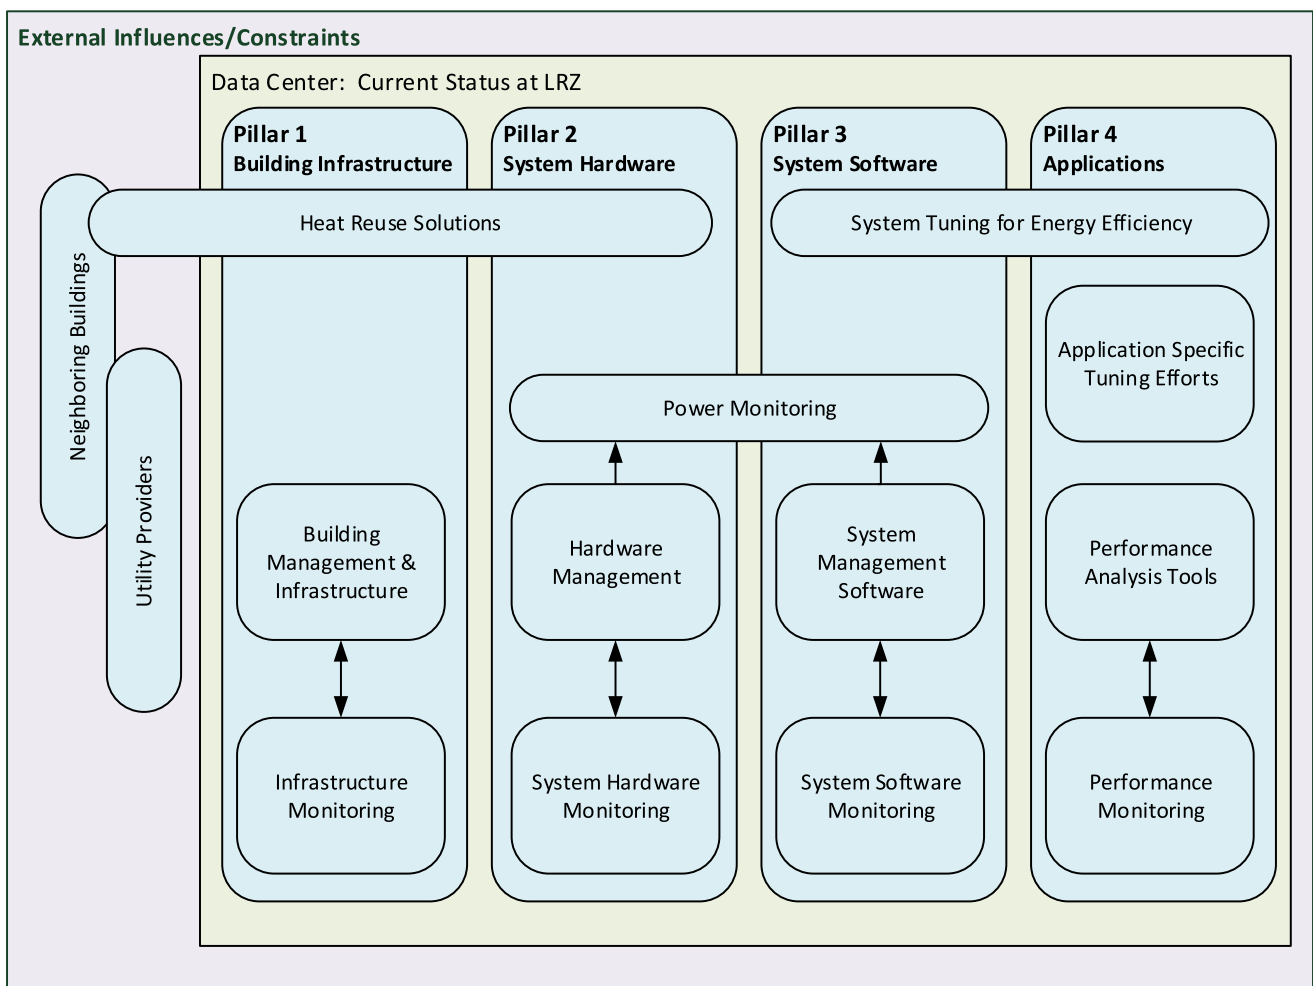
Fig. 3 High level energy efficiency activities at BAdW-LRZ

For heat reuse solutions BAdW-LRZ currently uses two different technologies. The waste heat from the SuperMUC supercomputer is used to heat a new office building whereas a smaller HPC system, the so called CooLMUC prototype, uses an adsorption chiller to indirectly cool an additional rack using energy from the system's waste heat.

In order to measure the energy consumed by a specific scientific application, a tool was needed that could collect power data from the HPC system hardware and correlate them with job information from the batch scheduling system. To address this, the PowerDAM V1.0 software was developed which collects and correlates data from pillars 2 and 3 [33].

For improved system tuning, DVFS support was enabled for the Slurm resource management system. Also, BAdW-LRZ is currently working with IBM to enable advanced energy aware scheduling techniques in the LoadLeveler resource manager. BAdW-LRZ is also actively involved in application performance improvements through participation in research projects such as AutoTune [3].

5 Mapping of future efforts at BAdW-LRZ to the 4 Pillar Framework

The 4 Pillar Framework for energy efficient HPC data centers can also be used to define and plan future activities. Figure 5 shows the needed developments and connections in and between each pillar in order to achieve the optimal energy efficiency for the BAdW-LRZ data center. Only a tight information exchange between all four pillars will enable tools that can provide substantial data center optimization to reduce the overall energy consumption. The aim is a data acquisition infrastructure that can collect important data from all 4 pillars as well as from the outside and the utility providers. This data can then be used by models and simulations to find optimized cross-pillar solutions that can be used by the infrastructure aware global resource manager to optimize the overall data center operation and to adjust the data center performance according to data center goals and policies.
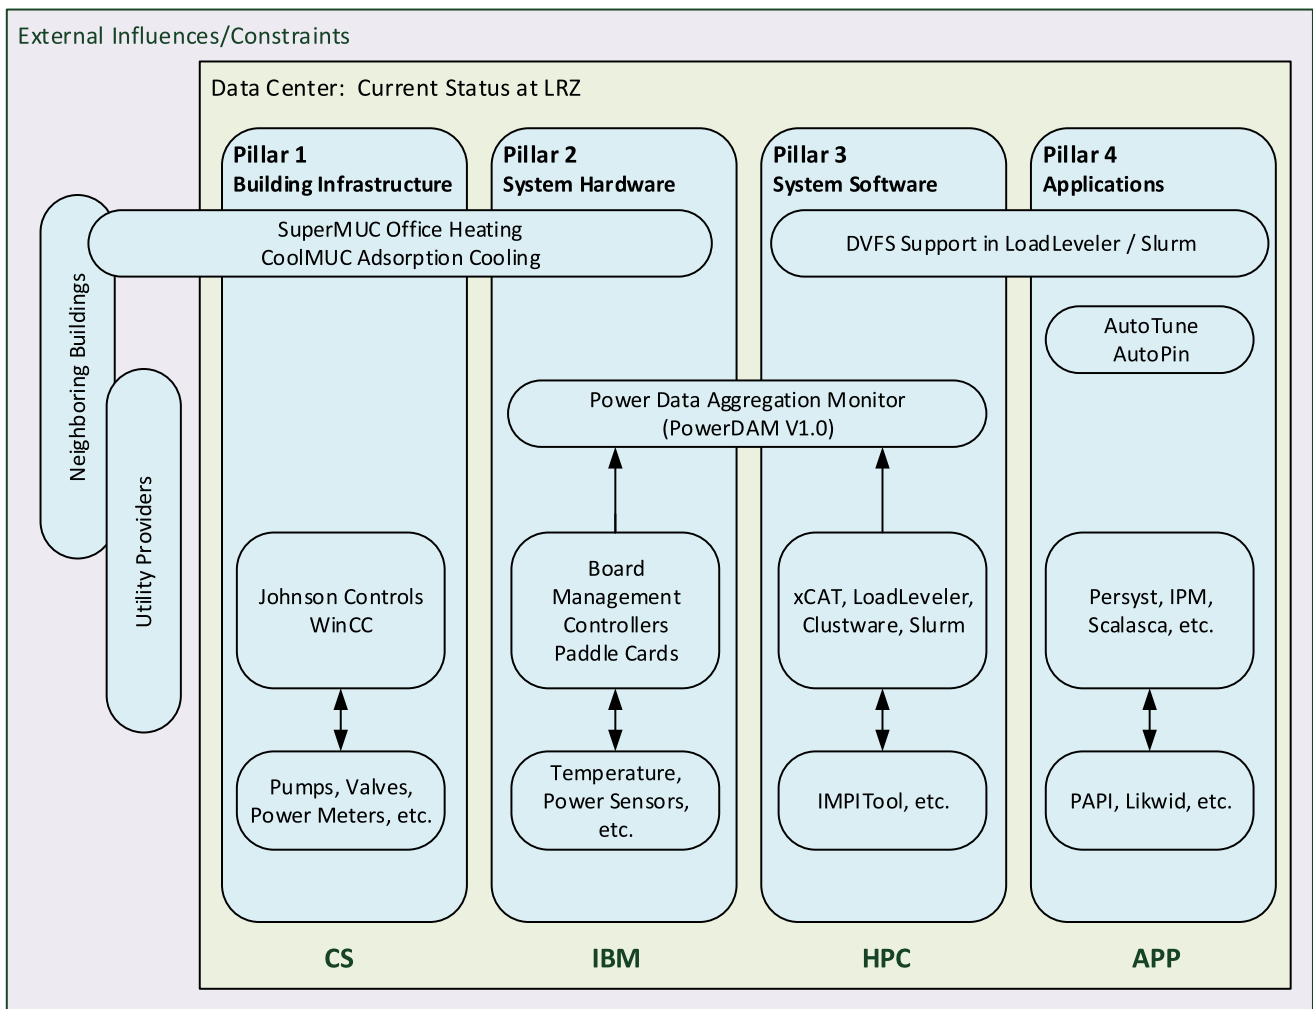
External Influences/Constraints

Data Center: Current Status at LRZ

Pillar 1 Building Infrastructure	Pillar 2 System Hardware	Pillar 3 System Software	Pillar 4 Applications

Neighboring Buildings

Utility Providers

SuperMUC Office Heating
CoolMUC Adsorption Cooling

DVFS Support in LoadLeveler / Slurm

AutoTune
AutoPin

Power Data Aggregation Monitor
(PowerDAM V1.0)

Johnson Controls
WinCC

Board
Management
Controllers
Paddle Cards

xCAT, LoadLeveler,
Clustware, Slurm

Persyst, IPM,
Scalasca, etc.

Pumps, Valves,
Power Meters, etc.

Temperature,
Power Sensors,
etc.

IMPI Tool, etc.

PAPI, Likwid, etc.

CS   **IBM**   **HPC**   **APP**

Fig. 4 Current energy efficiency activities at BAdW-LRZ

Table 1 Example project classification using the 4 pillar framework

Projects	Pillar 1	Pillar 2	Pillar 3	Pillar 4	Comments
CoolEMAll	covered (air cooled)	covered (air cooled custom hardware)	covered	not covered	
Fit4Green	not covered	not covered	covered (cloud based and service/enterprise portal oriented)	not covered	federation of data centers (multiple 4 pillar framework for each associated data center)
All4Green	covered (connects with utility provider)	not covered	not covered	not covered	builds on Fit4Green, data center federation (multiple 4 pillar framework for each associated data center)
GAMES	covered	covered	covered	not covered	cloud computing paradigm

6 Practical application

The 4 Pillar Framework can be used to classify current research projects. The following example shows a possi-
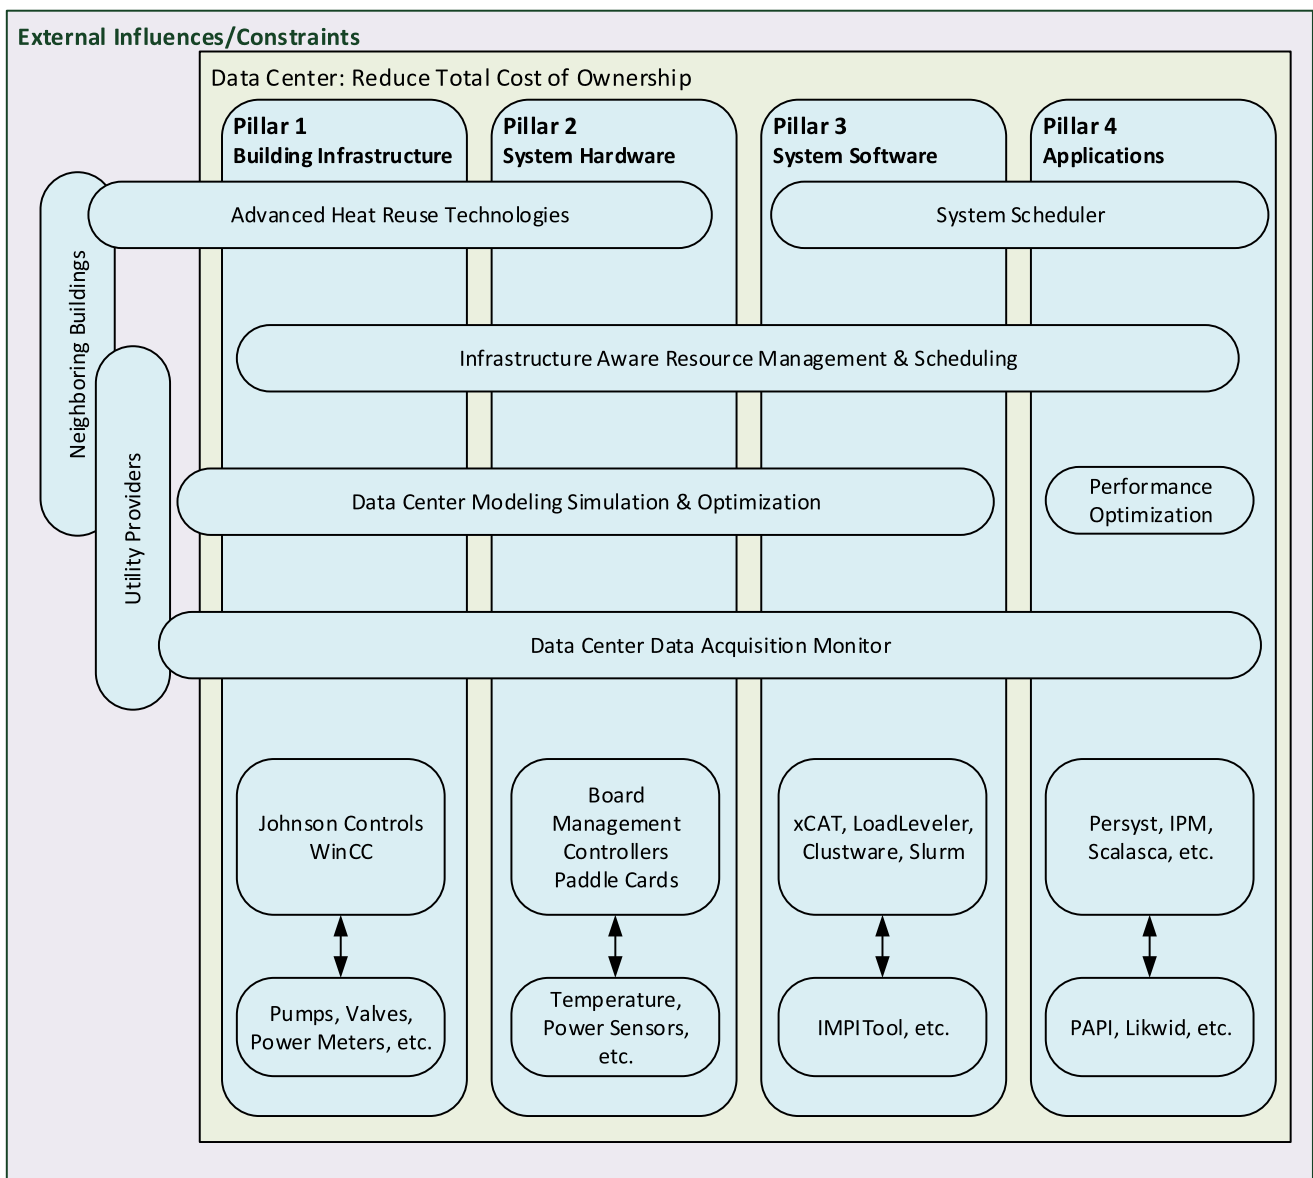
Fig. 5 Future energy efficiency plan at BAdW-LRZ

ble classifying for the Fit4Green project. The items below are taken from the 'Technical Approach' from the projects overview page [13].

1. "Fit4Green will introduce an optimization layer on top of existing data centre automation frameworks—integrating them as a modular "plug-in"—to guide the allocation decisions based on the optimized energy model based policies."—Usually the scheduling system makes the allocation decisions which is part of Pillar 3.

2. "Energy consumption models will be developed, and validated with real cases, for all ICT components in an IT solution chain, including the effects due to hosting data centres in sites with particular energy related characteristics, like alternative power availability and en-

ergy waste/recycle options, etc."—'all ICT components' could be interpreted as data center infrastructure and IT equipment which are part of Pillar 1 and 2, but because the emphasis is on the development of models and not tools there is no impact on pillar 1 and 2 from a practical (solution) point of view.

3. "Optimizations, based on policy modeling descriptions able to capture the variety of deployment that are possible for a given application or a set of applications, integrated with semantic attributes supporting the evaluation of energy consumption models, will guide the deployments decisions on which / when equipments needs to stay on, where/when applications should be deployed/relocated, also capitalizing on the intrinsic non linear behaviors of

energy consumption grow with respect to the load of ICT components."—'Applications' relate to Pillar 4, but the main goal of Fit4Green is a model for describing all possible deployments for an application. Therefore, this item will not impact Pillar 4 from a practical point of view. The deployment decisions relate to the scheduling system which is part of Pillar 3.

4. "Scenarios will be developed for several IT solutions deployed through there major categories of computing styles: Service/Enterprise Portal (a set of public services and/or private enterprise applications, typically accessible through Web interface, for a large variety of user agents); Grid (a supercomputer system accessible through the Grid middleware); and Cloud (a lab with cloud computing infrastructure based on open systems frameworks)"—This statement means that any solution developed by Fit4Green is only applicable to data centers that work in Service/Enterprise Portal, Grid, or Cloud computing style.

5. "There will be one pilot site for each computing style. Service/Enterprise Portal, Grid and Cloud pilots will support both single site and federated sites scenarios: multiple collaborating data centres inside the same organization for the Portal pilot; federation of supercomputer systems for Grid and open cloud federation of multiple labs for the Cloud."—The main target for the project seems to be a federation of data centers.

Using the information found in the 'Technical Approach' page, the Fit4Green project activities can be classified under Pillar 3 and are only applicable to a cloud, grid, or service portal based federation of data centers. This classification approach can be used on any research activity related to energy efficiency improvements for data centers.

Table 1 shows a possible classification for a couple of current research projects (CoolEMAll [7], All4Green [1], Fit4Green [12] and GAMES [14]). Column 2 to 4 (Pillar 1 to Pillar 4) shows which pillars are the main focus of the project (indicated by "covered") and any project specific assumptions or restrictions for each pillar. Column 5 "Comments" lists assumptions or restrictions related to the project in general.

This classification can help to find projects that are interesting for a specific data center. If, for example, one is looking for improvements in Pillar 1, one can remove Fit4Green from the list of project to look into. Similarly if one is not in a federation of data centers, one can focus on CoolEMAll and GAMES.

7 Conclusion

As demonstrated in this paper the "4 Pillar Framework for Energy Efficient HPC data centers" can be used by individual data centers, as well as research communities, as a foundation for all efforts related to the improvement of data center energy efficiency. It can be used to:

1. Identify data center specific areas of improvement
2. Identify required company resources and stakeholders
3. Guide current energy efficiency efforts
4. Plan future work
5. Present results and plans to stakeholders
6. Classify current research efforts

It is flexible enough to be a foundation for efforts related to:

1. Federation of data centers
2. Individual data centers
3. Individual data center parts and systems
4. Individual data center goals and policies

The 4 Pillar Framework has already shown its practical applicability and continues to play a significant role in the strategic planning of research activities and modernization efforts at BAdW-LRZ.

References

1. All4Green. http://www.all4green-project.eu/
2. ASHRAE. https://www.ashrae.org/
3. AutoTune. http://www.autotune-project.eu/
4. Bell RH Jr., Brochard L, Desota DR, Panda RD, Thomas F (2012) Energy-aware job scheduling for cluster environments. http://www.freepatentsonline.com/y2012/0216205.html
5. Browne S, Dongarra J, Garner N, London K, Mucci P (2000) A scalable cross-platform infrastructure for application performance tuning using hardware counters. In: Proceedings of the 2000 ACM/IEEE conference on supercomputing (CDROM), Supercomputing '00. IEEE Comput Soc, Washington. http://dl.acm.org/citation.cfm?id=370049.370424
6. Brunschwiler T, Meijer GI, Paredes S, Escher W, Michel B (2010) Direct waste heat utilization from liquid-cooled supercomputers. In: Proceedings of the 14th international heat transfer conference, Washington, p 23352
7. COOLEMALL. http://www.coolmall.eu/
8. Council J.N.B.I.P. (2012) DPPE: Holistic framework for data center energy efficiency—kpis for infrastructure, it equipment. operation (and renewable energy). http://home.jeita.or.jp/greenit-pc/topics/release/pdf/dppe_e_20120824.pdf

9. David H, Fallin C, Gorbatov E, Hanebutte UR, Mutlu O (2011) Memory power management via dynamic voltage/frequency scaling. In: Proceedings of the 8th ACM international conference on autonomic computing, ICAC '11. ACM, New York, pp 31–40. doi:10.1145/1998582.1998590

10. DEEP. http://www.deep-project.eu

11. Fan X, Weber WD, Barroso LA (2007) Power provisioning for a warehouse-sized computer. In: Proceedings of the 34th annual international symposium on computer architecture, ISCA '07. ACM, New York, pp 13–23. doi:10.1145/1250662.1250665

12. Fit4Green. http://www.fit4green.eu/

13. Fit4Green. http://www.fit4green.eu/content/overview

14. GAEMS. http://www.green-datacenters.eu/

15. Ge R, Feng X, Cameron KW (2005) Performance-constrained distributed dvs scheduling for scientific applications on power-aware clusters. In: Proceedings of the 2005 ACM/IEEE conference on supercomputing, SC '05. IEEE Comput Soc, Washington, p 34. doi:10.1109/SC.2005.57

16. The Green Grid. http://www.thegreengrid.org/

17. Energy Efficient HPC Working Group. http://eehpcwg.lbl.gov/

18. Hoefler T (2010) Software and hardware techniques for power-efficient hpc networking. Comput Sci Eng 12(6):30–37. doi:10.1109/MCSE.2010.96

19. Hsu CH, Feng WC (2005) A power-aware run-time system for high-performance computing. In: Proceedings of the 2005 ACM/IEEE conference on supercomputing, SC '05. IEEE Comput Soc, Washington, p 1. doi:10.1109/SC.2005.3

20. Fraunhofer IZM in Kooperation mit Fraunhofer ISI (2009) Abschaetzung des energiebedarfes der weiteren entwicklung der informationsgesellschaft. http://www.bmwi.de/Dateien/BMWi/PDF/abschaetzung-des-energiebedarfs-der-weiteren-entwicklung-der-informationsgesellschaft,property=pdf,bereich=bmwi,sprache=de,rwb=true.pdf

21. Johnsson (KTH) L, Netzer (KTH) G, Boyer (CINES) E, Graf (Juelich) S, Homberg (Juelich) W, Koutsou (CaSToRC) G, Jakic (IPB) J, Januszewski (PSNC) R, Puzovic (BSC) N, Roeblitz (SIGMA/UiO) T, Saastad (SIGMA/UiO) OW, Schembera (HLRS) B, Schwarz (Juelich) G, Shoukourian (LRZ) H, Strumpen (JKU) V, Thiell (CEA) S, de Verdière (CEA) GC, Wilde (LRZ) T (2013) PRACE 1IP-WP9 D9.3.3 report on prototypes evaluation. Public deliverable, waiting for European Commission approval for publication. http://www.prace-project.eu/PRACE-1IP-Deliverables?lang=en

22. Kim J, Chou J, Rotem D (2011) Energy proportionality and performance in data parallel computing clusters. In: Proceedings of the 23rd international conference on scientific and statistical database management, SSDBM'11. Springer, Berlin, pp 414–431. http://dl.acm.org/citation.cfm?id=2032397.2032431

23. Kim N, Austin T, Baauw D, Mudge T, Flautner K, Hu J, Irwin M, Kandemir M, Narayanan V (2003) Leakage current: Moore's law meets static power. Computer 36(12):68–75. doi:10.1109/MC.2003.1250885

24. Kogge P (Sept. 28, 2008) Exascale computing study: Technology challenges in achieving exascale systems. Univ of Notre Dame, CSE Dept Tech Report TR-2008-13

25. Koomey JG (2008) Worldwide electricity used in data centers. http://iopscience.iop.org/1748-9326/3/3/034008/pdf/1748-9326_3_3_034008.pdf

26. Koomey JG (2011) Growth in data center electricity use 2005 to 2010. http://www.twosides.us/content/rspdf_218.pdf

27. Lively C, Wu X, Taylor V, Moore S, Chang HC, Cameron K (2011) Energy and performance characteristics of different parallel implementations of scientific applications on multicore systems. Int J High Perform Comput Appl 25(3):342–350. doi:10.1177/1094342011414749

28. Loper J, Parr S (2007) Alliance to save energy, energy efficiency in data centers: A new policy frontier. http://www.ase.org/resources/energy-efficiency-data-centers-new-policy-frontier

29. Meijer GI (2010) Cooling energy-hungry data centers. Science 328:318

30. MontBlanc. http://www.montblanc-project.eu/

31. Patil V, Chaudhary V (2011) Rack aware scheduling in hpc data centers: an energy conservation strategy. In: 2011 IEEE international symposium on parallel and distributed processing workshops and phd forum (IPDPSW), pp 814–821. doi:10.1109/IPDPS.2011.227

32. PRACE. http://www.prace-ri-eu/

33. Shoukourian H, Wilde T, Auweter A, Bode A, Piochacz P (2013) Towards a unified energy efficiency evaluation toolset: an approach and its implementation at Leibniz supercomputing centre (lrz). In: ICT4S 2013: proceedings of the first international conference on information and communication technologies for sustainability, February 14–16, 2013. http://dx.doi.org/10.3929/ethz-a-007337628

34. Song S, Ge R, Feng X, Cameron KW (2009) Energy profiling and analysis of the hpc challenge benchmarks. Int J High Perform Comput Appl 23(3):265–276. doi:10.1177/1094342009106193

35. Valentine B Introducing sandy bridge. https://www.cesga.es/gl/paginas/descargaDocumento/id/135

Torsten Wilde received a M.Sc. degree in "Parallel and Scientific computation" from the University of Liverpool, UK, in 2000, and a Dipl.-Ing. (FH) in "Technische Informatik" from the University of Applied Sciences in Berlin, Germany. He is a research associate at the Leibniz Supercomputing Centre, Germany, working on methods of improving the energy efficient of data centers, a team member of the energy efficiency task force at LRZ, and a member of the Energy Efficient High Performance Computing Working Group (EEHPC WG). He is a work package lead for Partnership for Advanced Computing in Europe (PRACE).

Axel Auweter received a degree in "Computer Science" from Technische Univeriät Müchen, Germany, and subsequently joined LRZ's high-performance computing division in 2010 as a research associate. He is responsible for LRZ's research activities on energy efficient HPC system design and operation. In this role, Axel currently also acts as the leader of the work package on energy efficiency in the DEEP (Dynamical Exascale Entry Platform) project. His background is in computer architecture, system level programming and operating systems.

Hayk Shoukourian received a B.Sc. degree in Informatics and Applied Mathematics from Yerevan State University, Armenia, in 2010, and a M.Sc. degree in Informatics from the Technical University of Munich, Germany, in 2012. He is a Ph.D. student at Technical University of Munich and a research associate at the Leibniz Supercomputing Centre, Germany. He is working in the area of energy efficient data centers with special focus on models of time efficient management for energy conservation in high-performance computing.