The Energy Demand of Data Centers

Gunnar Schomaker, Stefan Janacek and Daniel Schlitt

Abstract Data centers are the backbone of today's information technologies. With increasing usage of cloud services and web applications, the need for remote computing and storage will only grow. However, one has to consider that increasing numbers of server and storage systems also mean increases in energy consumption. The power demand is caused not only by the IT hardware, but is also due to the required infrastructure such as power supply and climatization. Therefore, choosing the most appropriate components as well as architectural designs and configurations regarding energy demand, availability, and performance is important. This chapter depicts influencing factors and current trends for these design choices and provides examples.

Keywords Data center · Energy demand · Metrics · Renewable energy · Virtualization · Cooling concepts · Power supply

1 A Short Introduction to Data Centers

In our daily lives, we increasingly access data located "somewhere else," at a place most of us do not have to care about. Usually, this place is a data center, closer or farther from the end user, completely opaque with its high intrinsic technical complexity and specific requirements. This chapter offers a basic explanation of data centers and especially how their power consumption is pieced together, why they need to consume a significant amount of energy and how that turns information processing into a cost-intensive task.

Most basically, a data center is a dedicated building or at least a dedicated, separated room, exclusively held available for the placement of IT hardware. This definition deliberately excludes small storage rooms that many (smaller)

G. Schomaker (✉) · S. Janacek · D. Schlitt
R&D Division Energy, OFFIS, Oldenburg, Germany
e-mail: schomaker@offis.de

L.M. Hilty and B. Aebischer (eds.), *ICT Innovations for Sustainability*,
Advances in Intelligent Systems and Computing 310,
DOI 10.1007/978-3-319-09228-7_6

companies use to hold a rack containing some servers. While some define data centers to include the IT hardware, others do not. This is due to the fact that data centers can be run with different business models and operating strategies. A company might build a data center but not place a single server they own in it. Instead, third parties may buy their own servers and send them to the data center for it to manage their hardware. The data center's operating company only rents the "building" including all infrastructure needed to operate the servers, but not the servers themselves. This business model is called a co-location data center. Other data center operators construct and run the data center completely, including all hardware and software, selling only their services to their customers. A popular business model that uses exactly this scheme is very well known and roughly summarized by the term "cloud." This makes use of the possibilities gained by the technology of virtualization, which is explained later in this chapter. Here we assume that a data center contains the building, the IT hardware, and all infrastructure needed to successfully and safely operate it.

One of the main reasons why data centers have become complex and cost-intensive structures is reliability. This entails two areas: security and availability. It is fairly easy to explain why everyone using remote services or accessing data should care about security: to prevent others from spying on their personal or business files. For a data center, this not only means securing its digital entrances through firewalls and intrusion detection systems, but also reducing to the greatest possible extent the possibility of a break-in, the impact of natural disasters or misuse of any hardware by unqualified personnel. To prevent these incidents, data center buildings often resemble high security zones protected by video surveillance, road barriers, alarm systems, isolating devices, and fire protection systems. The building grounds of a data center should be above the levels of near water sources and be save from flooding.

Availability means the minimization of blackouts and deficiencies of the services. These may be a result of fluctuations in the power system or even entire power blackouts, but also of hardware failures or similar faults that may occur anytime during operation. To prevent service downtime or hardware damage, data centers use an elaborate combination of redundant systems for almost all technological layers. Starting with the power supply, a data center may use external emergency power supplies fueled by diesel and pre-heated to ensure the immediate operation until the outage can be resolved. There are also internal uninterrupted power supplies (UPS) with extensive battery packs that are able to power the data center for at least several minutes to bridge the time until the emergency power supply is activated or to filter a deterioration of quality caused by the energy supplier. In the worst case, at least the provided time frame allows the controlled shut-down of hardware to avoid damage or data loss.

The level of safety and redundancy is described by the tier level of a data center and finally results in reducing annual failing hours. The higher this level, the more redundancy the data center uses and the better it can guarantee a significantly higher availability [1]. However, all these security devices and redundant components come at a price: besides the acquisition costs of additional devices, the operational costs further increase due to more devices being run in parallel with reduced operational efficiency or at least in standby, waiting to take over in a predefined emergency case.

Data centers can be classified as follows:

- **Small data centers** contain 100–500 servers and have power inputs of about 50 kW.
- **Medium sized data centers** may host 500–5,000 servers with a power input of 240 kW.
- **Big data centers** host more than 5,000 servers, reaching a power input of about 2.5 MW, or even higher.

In Germany, for example, there are only about 50 big data centers, whereas there are about 370 medium sized data centers and 1,750 small ones [2].

2 Energy in the Data Center

A data center consists of several classes of devices, which all need a certain amount of power. Mainly these are the servers, the air conditioning, emergency power supplies and UPS, storage devices, network devices such as switches and routers, power distribution and other infrastructural devices, i.e. lighting, alarm or monitoring systems. Figure 1 shows a power consumption breakdown of these device classes. In efficient data centers, most of the power should be consumed by servers and other IT hardware whereas the air conditioning should consume less power; nevertheless, this is one of the main areas besides the servers that needs a significant amount of energy.

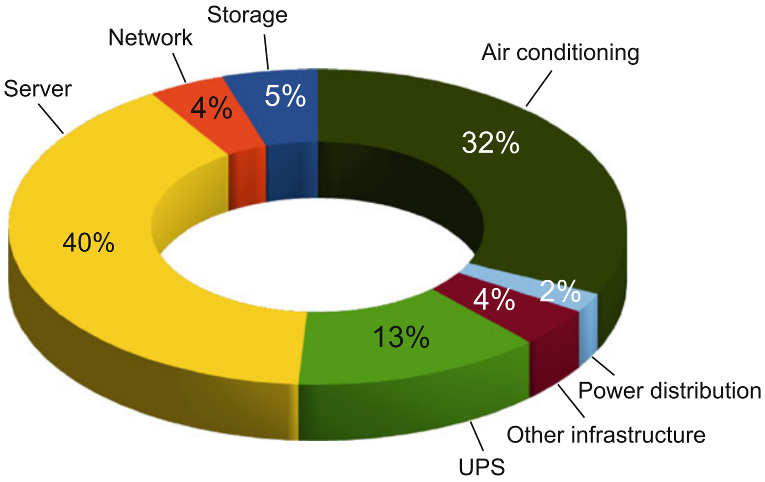

Fig. 1 Power breakdown of a typical data center; data from [3]

The air conditioning includes all components used to create cool air, distribute it and extract the hot air from the servers and devices. A common misconception is that only the servers generate waste heat and need to be cooled. It is true that the servers are the major heat producers, but each other device that consumes electric energy dispenses warm air, heating up the surrounding air and thus the entire room. If the electric components cannot be provided with enough cool air, serious hardware damage may result. Besides the servers, the UPS devices generate a significant amount of heat that needs to be transported out of the data center's room. Basically, according to the physical law of conservation of energy, the energy going into a system is equal to the energy coming out of the system. For a data center this means the energy in the form of power consumed must exit the data center, normally in the form of heat transported by hot air or water.

2.1 Dynamic Power Consumption and Efficiency Factors

When a new data center is designed and constructed or when old hardware gets replaced, the maximum power consumption of components is considered and all components are chosen wisely in order to keep within the maximum possible power a UPS might provide, for example. The relevant value for this process is mostly each component's plate power, thus the power printed on the back of each device, depicting its maximum theoretical consumption value. However, this value is seldom reached. First, the hardware manufacturers need to guarantee that their device will not extend its plate power, and as a result, the value is maximized. Second, modern power scaling technologies enable many devices to save power in times of low load, resulting in consumption values far below their plate power.

For servers, research has shown that a server's power consumption is significantly influenced by its current application load, mainly its processor's load [4]. Hence, if a server has less work to do, it will consume less power. Processor technologies such as clock/power gating or dynamic voltage and frequency scaling (DVFS), which changes the processor's voltage and frequency according to the current system load, have found widespread acceptance and lead to an even further improved dynamic power behavior of servers. This is also true of modern UPS and cooling devices, enabling the entire data center to adapt its power consumption to the current needs.

However, a fact that is often neglected is the efficiency factor of each device. The efficiency factor means that each device has a specific load range at which it shows the best relation of performance and consumed power. Operating a device out of this range results in a suboptimal efficiency, mostly implying a higher power consumption than it could have. The combination of several components in a complete system environment (e.g. a complete cooling chain) worsens this problem, as not every component can be operated at its optimal operation point. In this case, the best operation point holistically for all affected components should be found and used.

Considering these technologies and observations, the power consumption profile of a data center is mainly influenced by its application load profiles. If, for example, most of a data center's servers are used by employees of an office, and they follow regular working schedules, the power consumption profile of these servers may clearly show load peaks in the morning and afternoon, while being almost completely idle at nighttime.

One of the challenges of data center operators is to enable maximum performance operation during these peak hours while reducing the power consumption during low load times. Peak hours or low load, the accurate operation of all services must always be guaranteed. Later sections in this chapter show how this effort may be achieved by using technologies such as virtualization.

2.2 (Emergency) Power Supply

As stated above, the availability of the data center is a serious concern. Several redundant systems are thus used. The power supply system is often heavily redundant. To circumvent total power blackouts, on-site emergency power generators are used for a mid-term power supply. These are normally diesel-fueled generators standing near the data center's main facility. Even if the grid is stable and no power blackouts occur, these generators must be tested and refueled on a regular basis to prevent quality issues with older fuel.

The power path of a typical data center is shown in Fig. 2. Inside the data center, UPS systems are used to secure the power supply of the most important IT devices. These are the devices shown in the grey rectangle in the figure. Modern UPS systems are able to power these devices in case of power shortages for a few minutes. For this, they can fall back on a grid of battery packs. Similar to diesel generators, these need regular testing and maintaining. However, merely supplying power is not the only task for UPS systems in data centers. As many of the server and network devices are very sensitive to power or frequency fluctuations, the UPS also provides a netfilter functionality, dampening those effects and also preventing electrical surges.

Behind the UPS systems follows a wisely implemented power distribution system, consisting of modern power distribution units (PDU) that resemble normal multi-plugs, only with much more functionality. These units are able to measure and monitor the current power consumption of one rack or a single server; measured values are transmitted over ethernet connections to logging servers, managing this big data. Even remote shutdowns and switch-ons of single power ports are possible. These new possibilities enable the detailed energy logging of data center IT hardware and future automated power management of components lacking their own power saving functions. Of course, these intelligent, computerized PDUs also need power for themselves.

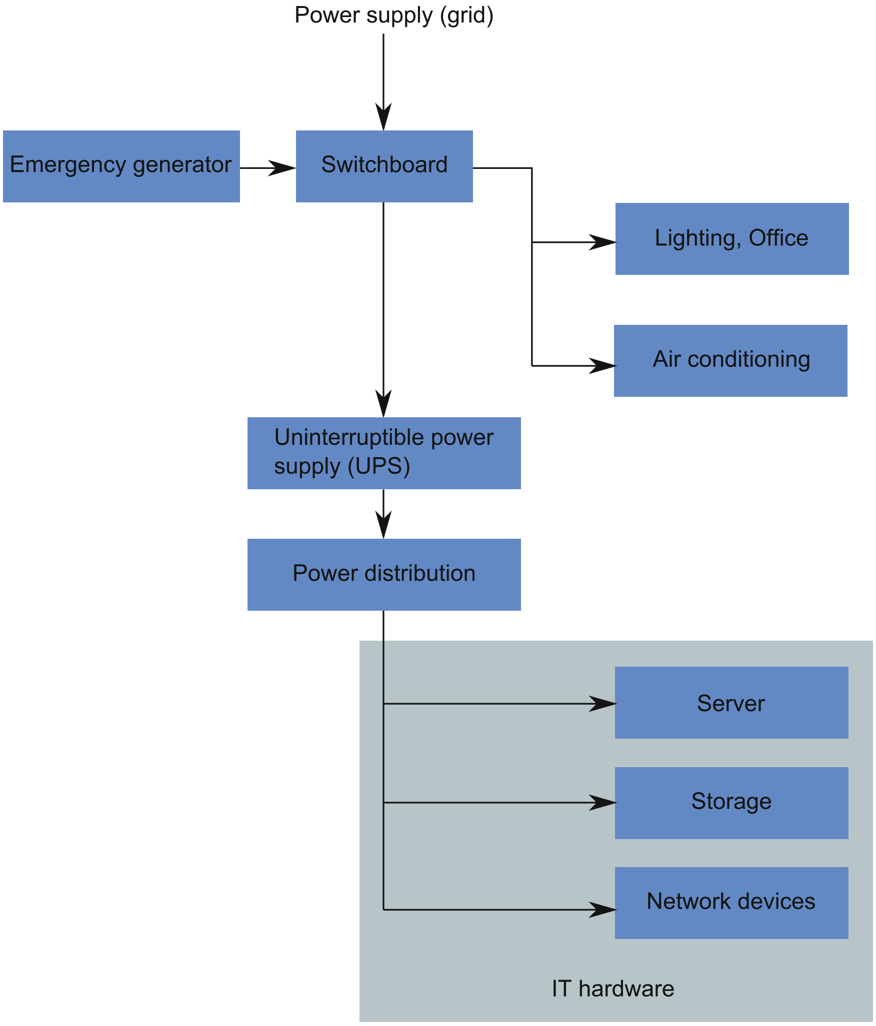

Fig. 2 Power path of internal data center components; IT hardware is highlighted in the *grey rectangle*

2.3 Hardware: Saving Potentials and Recent Developments

Over the last few years, the efficiency of data center hardware has changed considerably. A major trend has been to build faster and more compact servers, increasing their performance but also their energy density. This has lead to higher energy densities in racks, assuming that racks are not left half empty. Theoretically, a few of these modern servers are able to handle the work of many older ones, but because of rising workloads and requirements, the total number of servers has not decreased but rather often increased. As a result, new ways to handle this greater energy demand had to be found, specifically for both energy forms: incoming power and outgoing heat. Building thicker power tracks and higher dimensioned switchboards is relatively easy compared to the challenge of removing the heat from such a densely equipped rack in an efficient way. One possible solution is using cold aisle containments in combination with raised floor cooling. For a cold aisle containment, an encasement is built around two neighboring rack rows. Cold air is blown through the raised floor into this containment, and the servers in both racks may now take in this air through their fans. On the back of each rack, the heat leaving the servers is dispensed into the air, and a computer room air handler (CRAH) extracts it. The CRAH then cools the air with chilled water from a second cooling circuit. This technology leads to increased air conditioning efficiency.

An alternative method for heat removal in the highest energy density racks is to bring the chilled water not only to a CRAH but directly to the racks. From there the fully enclosed racks can be cooled by using in-row cooling devices, with air cooling devices in the raised floor, or by backdoor cooling. The purpose of in-row cooling and raised floor cooling devices is to minimize the distance the chilled air has to cover. Thus, the air loses less “cooling energy.” Backdoor cooling devices aim to eliminate heat directly at the server outputs in order to maintain a homogeneous room air temperature. These techniques are more complex to build as the chilled water circuit has to reach every rack. However, the higher cooling efficiency and thus the ability to increase the IT packing density are key reasons they will continue to be used.

The type of cooling technology and the desired cool air temperature both affect energy costs. The lower the desired temperature, the more energy will be spent to cool the air down to this temperature. In order to save energy for air conditioning, a possibility is to slowly increase the cool air temperature by a few degrees. The ASHRAE (formerly: American Society of Heating, Refrigerating and Air Conditioning Engineers) [5] offers standards and recommendations for data centers regarding this temperature value. However, the precise temperature depends on the used hardware, infrastructure and architectural design of the data center and of all components used. There are data centers that operate with a temperature in a range of 20–27 °C while others operate at up to 40 °C.

Besides these data center internal cooling technology trends, another question is how to efficiently cool down the heated water coming from the CRAH or the rack-based cooling devices. Conservative solutions use traditional refrigeration; however, this may be soon the most energy-inefficient way of cooling down the water. Instead, new solutions take advantage of natural conditions prevailing at the data center’s location. The simplest example of such a condition is a particularly low ambient temperature, enabling the data center’s cooling to apply free air cooling. Also, cold water reservoirs may be used.

2.4 Software: Saving Potentials and Recent Developments

Although the data center is mainly defined by hardware devices, some software trends have still had a major influence on the data center industry. One of them is the concept of virtualization. It allows the separation of physical servers and virtual machines that run the software provided to the users or customers. One physical server may execute several different virtual machines at the same time these share the physical machine's resources. This technology was originally introduced to improve the maintainability and flexibility of servers in a data center. Here, virtualization allows one to migrate virtual machines to a different server without stopping its running services. Old or faulty servers may be emptied in this way and then get substituted by new ones. However, this technology had much more potential than originally thought. The technology of live migration allows the migration of virtual machines between servers with practically no service downtime (in millisecond range) and within small time ranges of a few seconds. A recent trend is to use this concept to move as many virtual machines as possible to a single server and switch off currently not needed, “empty” servers [6, 7]. This allows significant energy savings, but also bears dangers. The data center's operator must guarantee to a certain extent that the execution of its services will proceed without problems. High variations in application load profiles may however lead to resource shortness on a physical machine when there are too many demanding applications. In this case, an early switch-on of more servers is needed to migrate some virtual machines to them, just enough to allow the accurate operation of all applications. This concept is known as load management and is currently the subject of heavy research and development, with some solutions already on the market [8].

2.5 Hardware Life-Cycle

Regarding the aspect of sustainability, life-cycles of data center hardware, especially servers, may become a concern. Normally, server hardware gets exchanged every few years, because faster and more energy-efficient hardware becomes available. Some of the discharged hardware may still be used for less important services, but at some point, this hardware will be disposed of. Hard disks will never be reused, since no company can take the risk of spreading stored data; as a result, these will be destroyed. Many companies exist that are dedicated to IT hardware recycling; however, the number of these devices that actually get recycled is hard to determine and reliable numbers are hard to find. Since UPS systems need regular tests and the life of batteries is limited, a similar problem arises with these. At least, many UPS distributors companies offer take-back of the batteries.

2.6 Energy Efficiency Rating

Most well-known energy efficiency metrics for data centers are based on the electrical energy as the sole analyzed energy form. The main differences between the metrics are the coverage of observed influencing parameters and thus their suitability for drawing certain conclusions. Some metrics only consider the actual used energy in a data center and lack the possibility to compare results between two data centers or within the same data center with different configurations. Other metrics try to relate the useful work to used energy, but fail to define a general concept for evaluating useful work.

The prevailing energy efficiency metric for data centers is the Power Usage Effectiveness (PUE) and its reciprocal Data Center Infrastructure Efficiency (DCiE). Both metrics were developed by The Green Grid [9]. The PUE is defined by total facility energy demand divided by IT energy demand with measurements over a whole year. However, its common application as a general energy efficiency metric for data centers is not quite correct. By definition PUE represents the additional energetical overhead of infrastructure components to run the IT systems. It is a good measure for evaluating optimizations on the infrastructure side, but once IT systems have been changed, the comparability is lost. As PUE is defined for a whole data center facility only, the partial PUE (pPUE) [10] has been derived from it to assess data center subareas.

A shortcoming of PUE is its inability to represent power dynamics in a data center. Thus, additional metrics have been proposed that focus on the dynamic power behavior of IT and infrastructure. PUE Scalability [10] by The Green Grid and Infrastructure Power Adaptability by Schlitt et al. [11] indicate the IT and facility/infrastructure power relation to rate the adaptability of infrastructure power in addition to the absolute overhead given by PUE.

In addition to PUE, there are several energy efficiency metrics with a focus on computing that can rather be applies as a general metric. These metrics assess power/energy demands in relation to the useful work done. As a whole, they only differ in their approach to how to assess useful work. However, all of them possess a subjective component, as the productive outcome (e.g. processed orders per time) of data center applications must be defined by humans. Thus, an application of such a metric is complex and unique for each data center. This effectively precludes a fair comparison between different data centers. Illustrative metrics are data center performance per watt (DCPpW) [12] by Dell and data center energy productivity (DCeP) [13] by The Green Grid. Because of the mentioned definitional problems, there are also eight proxy measures, which can be used instead of useful work. These proxies reduce the useful work essentially to performance or utilization.

For a more detailed insight into the energy efficiency of single IT hardware components, there are several kinds of energy benchmarks. Three well-known examples have been specified. (1) SPECpower_ssj2008 [14] by the Standard Performance Evaluation Corporation (SPEC) runs a server-side java application and measures the power demand of servers at 11 throughput levels. The benchmark delivers a performance/power value for each load level as well as an average value. (2) The TPC-Energy [15] benchmark by the Transaction Processing Performance Council (TPC) rates the energy efficiency of a full system under test conditions consisting of several server and storage systems with an interconnecting network. It runs an online transaction processing workload and measures throughput and power at full load as well as in idle. (3) SPC-1/E [15] by the Storage Performance Council (SPC) stresses storage (sub)systems with typical functions of business-critical applications and measures maximal throughput as well as the power demand at up to five load levels. Although these energy benchmarks describe the energy efficiency of IT components reliably, a high-level view of the facility is missing. However, in future energy efficiency metrics, system-level energy benchmarks will play a key role. An example of such a metric is the load-dependent energy efficiency (LDEE) [16], which is currently in development.

If the focus of data center assessment evolves from energy to resource efficiency, there are some other known metrics. The Energy Re-Use Efficiency (ERE) [17] modifies the PUE regarding the reuse of waste energy. If for example the data center's waste heat is used to heat nearby offices, that fraction of energy may be subtracted from the facility consumption in the PUE equation. Thus, ERE demonstrates the commitment to sustainability. The Carbon Usage Effectiveness (CUE) [18] represents the sustainability of data centers by relating the PUE with the carbon emissions produced. High carbon emissions per kWh result in a worse efficiency whereas low carbon emissions, for example by using renewable energy, may compensate for bad PUE ratings. The Water Usage Effectiveness (WUE) [19] takes the same line except it focuses on water usage instead of carbon emissions.

3 Renewable Energy and Energy Reuse

One of the most recent trends regarding sustainability and green data center operation is the emergence of the usage of renewable energy for data centers [20–23]. At first glance, this idea seems to contradict all the requirements of a data center: reliability and availability of services under almost all circumstances. Most of the currently available renewable energy sources are heavily intermittent; for example, wind generators and photovoltaics depend on the current weather situation, with the latter clearly seasoning power generation at nighttime. And power storage opportunities are still far behind the current needs. Despite all these obstacles, some companies have already started building "green" data centers. Most of them use a combination of different technologies and strategies to ensure their reliable operation. The first step is to construct on-site renewable energy generators such as wind generators, photovoltaic or biogas. These should be connected to the public grid to enable the usage of excess power by nearby consumers such as other industrial facilities or private households. Since at some time grid power will need to be purchased, CO_2 neutrality can be achieved by buying certificates or power from remote renewable energy generator parks. An example of a company with a green data center is Apple with its data center in Maiden, North Carolina, USA [24].

According to the physical law of conservation of energy, all energy going into a data center must leave it in some form. From a physical point of view, a data center may be seen as an energy converter, consuming electric power and generating heat. To render today's data centers more sustainable, the usage of this waste heat is an increasing challenge [25] and not always easy to achieve. One of the major aspects of waste heat usage is the temperature potential that needs to be reached in order to make waste heat usage beneficial. As an example, one Swiss data center heats a swimming pool with its generated waste heat [26].

4 Conclusion

Although the data center industry is a market of hard requirements, it still looked out for new energy-efficient and sustainable technologies. Besides the efforts data center operators and hardware distributors have made in this area, scientific research and development is ongoing, creating new and interesting ideas and products for future data center designs, architectures and more efficient components. One example of a relatively new architecture innovation is the switch towards a direct current (DC) power distribution in data centers, simplifying the power supply of servers and other devices. However, it is important to note that not all research efforts will lead to successful improvements and not every new trend will be adopted by the data center industry.

References

1. APC: Guidelines for specifying data center criticality/tier levels (2007)
2. Hintemann, R., Fichter, K., Stobbe, L.: Materialbestand der Rechenzentren in Deutschland. http://www.uba.de/uba-info-medien/4037.html (2010)
3. Hintemann, R.: Die Trends in Rechenzentren bis 2015, Ergebnisse einer neuen Studie im Auftrag des Umweltbundesamtes (2010)
4. Janacek, S., Schroeder, K., Schomaker, G., Nebel, W., Rueuschen, M., Pistoor, G.: Modeling and approaching a cost transparent, specific data center power consumption. In: 2012 International Conference on Energy Aware Computing, pp. 68–73 (2012)
5. ASHRAE: https://www.ashrae.org (2014)
6. Hoyer, M., Schroeder, K., Schlitt, D., Nebel, W.: Proactive dynamic resource management in virtualized data centers. In: Proceedings of the 2nd International Conference on Energy-Efficient Computing and Networking (2011)
7. Tantar, A.A., Tantar, E., Bouvry, P.: Load balancing for sustainable ict. In: Proceedings of the 13th Annual Conference Companion on Genetic and Evolutionary Computation, GECCO '11, pp. 733–738. ACM, New York (2011)

8. VMware: Distributed resource scheduler und distributed power management. http://www.vmware.com/de/products/vsphere/features-drs-dpm (2014)
9. Belady, C., Rawson, A., Pflueger, J., Cader, T.: The Green Grid Data Center Power Efficiency Metrics: PUE and DCiE (2008)
10. Avelar, V., Azevedo, D., French, A.: PUE: A Comprehensive Examination of the Metric. White Paper (2012)
11. Schlitt, D., Schomaker, G., Nebel, W.: Gain more from PUE: assessing data center infrastructure power adaptability. In: 3rd International Workshop on Energy-Efficient Data Centres (2014)
12. Pflueger, J.: Re-defining the 'green' data center (2008)
13. Haas, J., Monroe, M., Pflueger, J., Pouchet, J., Snelling, P., Rawson, A., Rawson, F.: Proxy proposals for measuring data center productivity. Consortium Green Grid (2009)
14. Lange, K.: Identifying shades of green: The SPECpower benchmarks. Computer **42**(3), 95–97 (2009)
15. Poess, M., Nambiar, R., Vaid, K., Stephens Jr, J., Huppler, K., Haines, E.: Energy benchmarks: a detailed analysis. In: Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking, pp. 131–140. ACM (2010)
16. Schlitt, D., Nebel, W.: Load dependent data center energy efficiency metric based on component models. In: Proceedings of 2012 International Conference on Energy Aware Computing (ICEAC 2012), pp. 20–25. IEEE (2012)
17. Patterson, M., Tschudi, B., Vangeet, O., Cooley, J., Azevedo, D.: ERE: A metric for measuring the benefit of reuse energy from a data center. The Green Grid: Metrics and Measurements. White Paper (2010)
18. Belady, C., Azevedo, D., Patterson, M., Pouchet, J., Tipley, R.: Carbon Usage Effectiveness (CUE): A Green Grid Data Center Sustainability Metric. White Paper, vol. 32 (2010)
19. Patterson, M., Azevedo, D., Belady, C., Pouchet, J.: Water Usage Effectiveness (WUE): A Green Grid Data Center Sustainability Metric. White Paper, vol. 35 (2011)
20. Ren, C., Wang, D., Urgaonkar, B., Sivasubramaniam, A.: Carbon-aware energy capacity planning for datacenters. In: Proceedings of 2012 IEEE 20th International Symposium on Modeling, Analysis Simulation of Computer and Telecommunication Systems (MASCOTS), pp. 391–400 (2012)
21. Ghamkhari, M., Mohsenian-Rad, H.: Optimal integration of renewable energy resources in data centers with behind-the-meter renewable generator. In: IEEE International Conference on Communications (ICC) (2012)
22. Zhang, Y., Wang, Y., Wang, X.: Greenware: greening cloud-scale data centers to maximize the use of renewable energy. In: Kon, F., Kermarrec, A.M. (eds.) Middleware 2011, Lecture Notes in Computer Science, vol. 7049, pp. 143–164. Springer, Berlin (2011)
23. Mohsenian-Rad, H., Leon-Garcia, A.: Coordination of cloud computing and smart power grids. In: Proceedings of IEEE Smart Grid Communications Conference, Gaithersburg, Maryland (2010)
24. Apple: Apple und der umweltschutz. http://www.apple.com/de/environment/progress/ (2014)
25. Brunschwiler, T., Smith, B., Ruetsche, E., Michel, B.: Toward zero-emission data centers through direct reuse of thermal energy. IBM J. Res. Dev. **53**(3), 11:1–11:13 (2009)
26. Brodkin, J.: Swimming pool heated by data center's excess heat (Computerworld) (2008)