[·] Andreas Berl and Hermann De Meer. A virtualized energy-efficient office environment. In *Proc. of the 1st ACM SIGCOMM Int'l Conf. On Energy-Efficient Computing and Networking (e-Energy 2010)*, pages 11–20. ACM, 2010.

See `http://www.net.fim.uni-passau.de/papers/Berl2010c` for full reference details (BibTeX, XML).

A Virtualized Energy-Efficient Office Environment

Andreas Berl
University of Passau
Innstr. 43
Passau, Germany
berl@uni-passau.de

Hermann de Meer
University of Passau
Innstr. 43
Passau, Germany
demeer@uni-passau.de

ABSTRACT

The energy efficiency of information and communication technology becomes more and more important due to the raise of energy costs and the world wide desire to reduce CO_2 emissions. Data centers have been in the focus concerning their energy efficiency lately, however, also office hosts that are located outside of data centres consume huge amounts of energy (e.g., in public administration or companies). Such office environments promise a high potential in terms of energy savings – a significant number of hosts remains to be turned on 24 hours per day while being mostly underutilized. This paper investigates the energy consumption in office environments and discusses the potential of energy savings. An energy-efficient office management approach is suggested, based on resource virtualization, power management, and resource sharing. Additionally, the paper evaluates simulation results concerning energy consumption and service provision in the managed office environment. The evaluation indicates that about 75% of energy savings are achievable in office environments without a significant interruption of provided services.

Categories and Subject Descriptors

C.2.4 [**Distributed Systems**]: Distributed applications

General Terms

Energy efficiency, office environment, virtualization, peer-to-peer, power management, resource sharing

1. INTRODUCTION

Energy efficiency of information and communication technology has become an important topic in companies and public administration – the bottleneck of costs has changed. While hardware costs are decreasing on the one hand, costs of energy are increasing on the other hand. In addition, there are world wide efforts to turn IT green, (e.g., CO_2 emissions need to be reduced). Data centres are well known and often discussed consumers of energy. Koomey [13], e.g., reports that data centres in the USA and worldwide have doubled their energy consumption from 2000 to 2005. However, also end-devices have considerably contributed to the increase of electricity consumption, according to a 2006 survey [5] commissioned by the EU.

Office hosts outside of data centres contribute significantly to the overall IT energy consumption, simply because of the high number of such devices – in offices usually each employee has his own host. It is important to see that office hosts do not only consume energy while users are sitting in front of them. Instead, hosts are often running on a 24/7 basis. Such hosts are running due to several reasons: Users process overnight jobs (e.g., downloads, backups, or test runs), users need remote access to their hosts from outside the office (e.g., users work from home or are currently at a customers office), or users simply forget to turn off their hosts when they leave the office. Even when hosts are in use, they are often underutilized (e.g., in terms of CPU load) by typical office applications. Unfortunately, unused and underutilized hosts still consume a considerable amount of energy (see Section 2).

Several approaches have been suggested that deal with high energy consumptions of hosts in office environments (see Section 5). Such solutions range from the enforcement of office-wide power-management policies to thin-client approaches, where users share resources on terminal servers. As extension to power-management solutions and opposed to data-centre based terminal-server approaches, this paper suggests to combine an office-wide power management with distribu-ted resource sharing in office environments. It presents a managed office environment based on virtualization methods that performs a shift from the currently available distributed local resource management (per user) towards a centralized global resource management (per office). The number of simultaneously running hosts in the office environment is reduced, while the utilization of the hosts is raised. This enables a major reduction of the overall energy consumption within the office, without significantly decreasing quality or quantity of provided services.

The remainder of this paper is structured as follows: Section 2 investigates the energy consumption of office environments and identifies the potential of energy savings. Section 3 describes general methods to achieve energy-efficiency in office environments and points out requirements that have to be met in the development of a managed office environment architecture. Additionally, the section discusses the virtualization of office resources and suggests a managed office environment architecture. Section 4 compares energy consumption and provided services of the suggested managed office environment with a common office environment. It illustrates that the suggested architecture has the ability to save up to 75% of the energy costs of today's offices, while providing similar services to the users. Section 5 discusses related approaches to achieve energy efficiency in office environments and Section 6 concludes this paper.

2. POTENTIAL OF ENERGY SAVINGS

Hosts in office environments are often running without being physically used. This happens for short time periods, e.g, if users are in meetings, do telephone calls, have lunch or coffee breaks, etc. But it also happens for longer periods of time. Physically unused hosts are often turned on, because users (or administrators) access them remotely. Remote access happens from the user's home (e.g., in the evenings, or on home working days) or when users are working externally (e.g., at a customer office). Remote access is needed in such cases to access office specific applications and data. The user may need access to email accounts, personal data (e.g., documents, addresses of customers, data in data bases), or applications (e.g., special office/graphics/database applications). Another important cause that leads to physically unused but running hosts is overnight jobs. A user (or administrator) might schedule a certain job for the night, e.g., a simulation, a download, a backup, or the defragmentation of the hard disk. It is intended to run the job during the time the user is not using his host, so the job does not interfere with usual work. Apart from such reasons, some users simply forget to turn off their hosts, when they leave the office. Webber et al. [18] have analyzed sixteen sites in the USA and reported that 64% of all investigated office hosts where running during nights.

Hosts that are idle (0% CPU load) still consume a considerable amount of energy, compared to computers that are turned off. Different hardware components of a host need to be supplied with power even when the host is idle. Table 1 illustrates the relation of the energy consumption of loaded and unloaded hosts. It refers to measurements that have been performed at the University of Sheffield on hosts that are typically used as personal computers [6]. The four columns of Table 1 illustrate the type of the host (Host), the energy consumption of the hosts when they are turned off (Standby), the energy consumption when they are unloaded (Idle), and the energy consumption when they are intensively used (Intense). Intensively used means in this case, that the host is actively doing arithmetic on a large data set and writing results back to disk. A highly important fact is that idle hosts still consume 49% to 78% of the energy that they need in the intense usage scenario, which contributes to a big part of the overall energy consumption. Most current hosts provide low-power modes that can be configured by the user and kick in when a host is idle. Power management strategies include, e.g., slowing down a host's clock rate, turning off power to certain circuits, powering down the hard drive, powering down the monitor, or hibernating the complete host. Therefore, low-power and hibernation modes can save a significant amount of energy, compared to the idle state. However, as a matter of fact, many devices that are low-power capable do not successfully enter such modes. Low-power modes are subject to the complex combined effects of hardware, operating systems, drivers, applications, – and after all – the user-based power management configuration. Webber et al. [18] report that in the investigated offices only 4% of all hosts actually have switched to low-power modes during the night.

Host	**Standby**	**Idle**	**Intense**
Dell Optiplex SX280	1	58 (58%)	100
HP dx5150S	2	43 (49%)	87
Macintosh Mini	2	22 (59%)	37
Viglen VM4 Cube	3	61 (56%)	108
Viglen Genie	1	92 (62%)	149
Viglen EQ100	3	46 (78%)	59
Dell Optiplex 210L	2	70 (52%)	135
Viglen Genie Core Duo	2	65 (76%)	86
IBM X40 Portable	2	27 (73%)	37
Average	2	54 (63%)	89

Table 1: Energy consumption in watts – personal computers at the University of Sheffield [6]

Apart from idle hosts, also loaded hosts aren't used in energy efficient ways in office environments. Hosts are usually underutilized by typical office applications (e.g., text processors, browsers, or mail clients). This leads to a high number of lightly utilized hosts that consume nearly as much energy as heavily utilized hosts. Additionally, common office environments do not distinguish between local and remote usage of hosts. Users that are physically working on hosts are equally served as users that work remotely (e.g., via VNC software[1]). However, only users with physical host access need a separate host to work with. Remote users – and also jobs without any user interaction (e.g., backups, downloads, etc.) – do not necessarily need to utilize separate hosts. Local users could share their resources with remote users to increase the utilization of hosts in the office environment.

In energy efficient office environments, idle hosts should consume considerably less energy than loaded hosts, independent of the user's power management configuration and independent of broken low-power modes. In addition, the energy consumption of loaded hosts should be limited. This can be done by enabling a resource sharing that raises the utilization of hosts and reduces the number of used hosts in the office.

3. A MANAGED OFFICE ENVIRONMENT

This section suggests a managed office environment that exploits the potential of energy savings (as it is described in Section 2) by virtualizing host resources and managing them in energy efficient ways.

3.1 Methods and Requirements

When a user powers on his host in a common office, he finds his usual working environment. This working environment is called *personal desktop environment (PDE)* in this paper and typically consists of an operating system, applications, and the user's personal configurations. Although, in common offices often roaming profiles are available (see Section 3.2), the PDE as a whole is fixed, i.e., it is bound to a certain host in the office. When the PDE is turned on/off, also the host is turned on/off and vice versa. Users are able

---

[1] http://www.realvnc.com/

to access their PDE locally within the office or they may also be able to access it remotely from outside the office.

In the managed office environment, PDEs are additionally used as *mobile services*. Mobile services are freely movable within the office environment and are used to achieve service consolidation. When the user is not physically using his office host, his PDE can be decoupled from the host and be migrated to another host for energy reasons. Several PDEs can be provided by a single host. Therefore, a user's host is not necessarily turned on when a user utilizes his PDE – the PDE may be provided by a different host. In Figure 1 the transition from a common to a managed office environment (based on PDEs) is illustrated. It can be observed that in the common office environment the PDEs and the hosts are interdependent. Seven hosts are turned on in Figure 1 a) together with seven PDEs. Three hosts (with PDEs) are turned off. The situation is very different in the managed office environment Figure 1 b). Although the number of currently running PDEs is the same as in Figure 1 a), only four hosts are actually turned on. It can be observed, e.g., that the upper right host is providing three PDEs to users simultaneously. Based on the availability of mobile PDEs, energy efficiency is achieved in three steps:

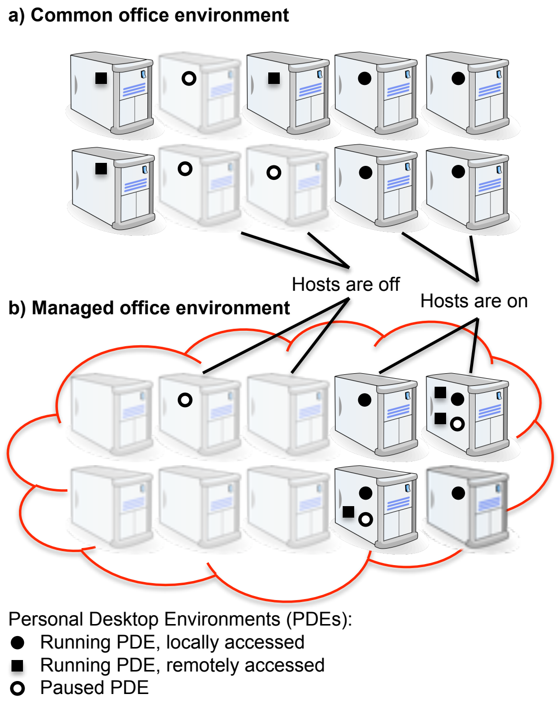

Figure 1: Common and managed office environment

- Unused PDEs in the office environment are stopped from consuming resources. If a PDE is idle (no job is performed on behalf of its user) it will be suspended.
- Used PDEs are consolidated on a small number of hosts. If a PDE is not accessed locally (the user does not physically access his office host), the PDE becomes a mobile service and may be migrated to other hosts to achieve consolidation.
- Hosts that do not provide running PDEs are shut down to save energy.

Suspending PDEs allows an optimization of the energy consumption E of each host h_i in the office environment, where $1 \leq i \leq n$ and n is the number of hosts. This approaches an energy consumption of $\sum_{i=1}^{n} min(E(h_i))$ within the office. The consolidation of PDEs allows an optimization of the overall energy consumption within the office environment, by considering the office as a whole, approaching an energy consumption of $min(\sum_{i=1}^{n} E(h_i))$.

To realize the envisioned energy-efficient management in office environments, several requirements have to be met. Hardware resource sharing (e.g., CPU cycles, memory, or disk space) among hosts in office environments is necessary in order to make idle resources available for PDE of other office hosts. A runtime environment has to be established, where PDEs of other users can be processed. In the best case, this runtime environment is flexible enough to enable the processing of a wide variety of different PDEs. PDEs might consist of different operating systems (e.g., Windows, MAC, or Linux) or even be executed on different computer architectures (e.g., x86 or PowerPC). A clear separation between different PDEs and the host they are executed on needs to be achieved, in order to prevent interference. Additionally mechanisms have to be applied that enable to power off unused hosts (to save energy) and to power them on again if they are needed again to provide additional services for users. PDEs need to be suspended and stopped from using resources, if they are idle. When the user wants to access the PDE again, it has to be resumed as fast as possible. Additionally, it has to be movable from one host to another, without terminating the processes that are currently running within the PDE. A temporary pause of process execution may be tolerated (similar to closing and opening a laptop), however, after that pause the PDE should continue to operate as expected by the user. All of the office hosts within the office environment have to be logically connected in order to enable a mediation of free resources and PDEs. In common offices this kind of interconnection is not available, but in the managed office it is necessary, because the states of PDEs and hosts are changing over time and PDEs may change their locations. The managed office environment needs a management entity that 1) suspends currently unloaded PDEs and 2) consolidates loaded PDEs on a small number of hosts and 3) powers down unused hosts. The consolidation process requires a reasonable and energy-efficient mapping (scheduling) of PDEs to hosts in the dynamic environment. It is important to see that the energy-efficient management can only take place under the precondition that services which are provided to users remain similar to usual office services in terms of quantity and quality. The energy-efficient operation of the managed office needs to be achieved, without significantly interrupting the day to day work of users. Minor changes in the usage of office hosts, however, may be tolerable by users. Summarized, the requirements of an energy-efficient management in office environments are defined as follows: 1) Hosts need to provide hardware resources for PDEs in separated runtime environments. 2) Hosts need to be shut down (or hibernated) and to be powered up again, if necessary. 3) PDEs need to be suspended when idle and to be resumed if necessary. 4) PDEs have to be movable from host to host (mobile services). 5) Addressing and mediation of office hosts (resources) and PDEs (services) has to be enabled. 6) An energy-efficient management of hosts and PDEs has to be performed, without interrupting the daily work of users.

3.2 Virtualization

A first important virtualization approach that is used in the managed office environment is system virtualization. It enables service consolidation and is successfully applied to data centres today. It can be adopted to office environments in order to achieve a similar utilization and energy efficiency of office resources. In system virtualization *virtual machines (VMs)* are created from idle resources. Full hosts are virtualized, consisting of virtual CPUs, virtual memory, virtual hard disk, virtual network interface card, etc. A VM is an imitation of a real machine in such a way that an operating system can be installed on it without being aware of the resource virtualization. The software that provides VMs is usually called *Virtual Machine Monitor (VMM)* (e.g., VMWare Server[2], QEMU [2], or Xen [1])) and is able to process several VMs simultaneously on a single host. There are several basic primitives of management functions available for VMs: create, destroy, start, stop, migrate, copy, pause, and resume VM. It is even possible to have a *live migration* [7]. This means that a service in a VM can be migrated to another host without being interrupted.

A PDE, as it is described in Section 3.1, can be encapsulated within a VM and inherits all of the VM-related features. Therefore, system virtualization, together with its management functions concerning VMs meets several of the requirements stated in Section 3.1. It enables the operation of PDEs in separated runtime environments (VMs). VMM can trigger the shut down of a host if required. Hosts can be powered up again, e.g. by using wake on LAN mechanisms[3], to boot into the VMM again. PDEs can be suspended by the VMM if they are idle and be resumed again if necessary. A major advantage of this kind of PDE suspension (compared to other management solutions, as discussed in Section 5) is that it works completely independent of the low-power modes and capabilities within the PDE. Additionally, when PDEs are enclosed in VMs they can be migrated from host to host, without a durable interruption of running services.

However, the costs of migration (as discussed in [3]) are a problem in the office environment. Whereas in data centres usually only processes are migrated (operating system and data are typically stored on network storage), PDEs have to be migrated entirely. This leads to considerable overhead because operating system and user data and applications might sum up to several GBs of data. To reduce this overhead, a *standard PDE (SPDE)* is stored on every office host in the managed office environment. The SPDE is a preconfigured full featured operating system (e.g., Windows or Linux), together with common applications, that provides the basis for each PDEs. Users can derive their own PDE from the SPDE (e.g., by installing additional applications or storing data). When an PDE is migrated from one host to another, not the complete PDE is transferred. Instead, the difference $DIFF = (PDE{-}SPDE)$ is migrated, consisting only of the user's personal changes. The receiver can recover the PDE from $DIFF$. Furthermore, if the PDE is re-migrated back to the original host, a second difference can be calculated that only contains current changes, further minimizing the network traffic. Additionally, the migration of PDEs can be supported by the application of roaming profiles within the office environment. Roaming profiles are often available in offices and enable a mobility of user profiles within the office (e.g., based on samba[4]). Users are able to log on to different machines within the office and access their personal software configuration using data centre-based network-storage solutions. This way, the data that needs to be migrated within the office is reduced and the performance of migrations is increased. Another comparable migration of PDEs from host to host has been discussed in the *Internet Suspend/Resume* project [14]. In this project, PDEs are migrated from desktop to desktop in order to enable pervasive personal computing. Distributed file systems (e.g, OpenAFS[5]) are used to reduce the amount of data that has to be transferred. A so called *transient thin-client mode* is under development, which allows the transient switching from remote PDE access (thin-client mode) to direct PDE access (thick-client mode) during its migration.

A second important virtualization approach that is needed to realize the managed office environment is based on P2P technology. Independent of the logical network that is used to interconnect hosts, the resource sharing in the managed office environment is done in P2P manner. There is no central element that provides resources to run PDEs on, as it is available in the thin-client/terminal-server approach. Instead all of the office hosts are sharing their resources. Therefore, methods of P2P overlays can be used to realize a *management environment (ME)* that interconnects hosts and provides mediation for hosts and PDEs. P2P content distribution networks (e.g., eDonkey[6] or BitTorrent[7]) are often used to share files among users. Such protocols provide several functions that can be adopted to office environments. First, these kind of networks create and maintain an overlay network among participants that enables a logical addressing of hosts, users, and content. Second, they enable the mediation of resources and are able to bring providers and consumers of content together. Third, such networks additionally manage the access to resources, in order to achieve an optimal and fair distribution of resources among all users of the network.

Concerning office environments, P2P overlays are able to meet several requirements as defined in Section 3.1. P2P overlays enable interconnection, addressing, and mediation of PDEs and hosts within the office environment. They also enable a management of PDEs and hosts based on their current states (e.g, powering off/on hosts or PDEs). Three different approaches of P2P overlays are possible to realize the ME in the office environment: 1) centralized client/server-based 2) pure P2P-based 3), and a hybrid approach. The most simple approach in terms of setup, administration, and management is the centralized client/server-based approach [15], where one or more dedicated servers are managing the office. All hosts of the office environment are logically connected to the centralized server and are reporting state changes. Although the ME is client/server-based, the resource-sharing is still done by office hosts in P2P manner. The peers of this P2P overlay are the hosts of the office environment – PDEs are not aware of the P2P net-

---

[2] http://www.vmware.com/de/products/server

[3] http://www.energystar.gov/index.cfm?c=power_mgt.pr_power_mgt_wol

[4] http://www.samba.org/

[5] http://www.openafs.org/

[6] http://www.overnet.org/

[7] http://www.bittorrent.com

work. The main downsides of this approach are the establishment of a single point of failure, and the fact that additional (energy consuming) hardware is needed to realize the management. A pure P2P-overlay approach (e.g., Chord-based [16]) realizes management in the office in a complete decentralized way – hosts are managing themselves. Each is additionally a management instance, no additional hardware is required and no single points of failure are added to the system. However, this approach adds complexity to the system. No global view on the office environment is available and the global goal of energy-efficiency has to be achieved by distributed management algorithms. A hybrid P2P overlay approach is a compromise between the centralized and the pure P2P approach. One or more centralized management entities are enclosed within VMs, similar to PDEs. This kind of distributed management enables the establishment of a managed office without imposing a need for further infrastructure elements. Available hardware can be used to provide the distributed terminal service in an energy-efficient way, however, the load within the managed office environment is increased by additional VMs. Similar to PDEs, the management entities are movable and can therefore be consolidated with other PDEs in order to save energy. The managed office environment in this paper is based on the centralized client/server ME approach (in a first step). In future work, this approach will be extended to the described hybrid ME approach.

3.3 Architecture

This section describes an architecture that achieves energy-efficient management in office environments, based on resource virtualization. The architecture has been simulated and results are discussed later on in Section 4.

PDEs: A PDE is enclosed in a VM and consists of five parts: 1) An operating system, 2) applications, 3) personal user data, 4) the user's personal configurations, and 5) a small driver that enables a communication with the VMM. The communication with the VMM enables an extended monitoring of user behaviour and resource usage within the PDE. In the managed office environment, PDEs are decoupled from hosts and the PDE can change into different states, based on the users behaviour. PDE states are:

- FIXED – The user physically accesses his office host. His PDE can not be migrated to another host for consolidation.
- MOBILE – The user doesn't access his office host physically. This means either the user uses his PDE remotely or the PDE is processing a job without user interaction. The PDE is movable within the office.
- PAUSED – The user does currently not use his PDE, therefore it can be suspended. Suspended PDEs don't consume resources on any host, except of disc space.

The states of PDEs are illustrated in Figure 2. In a common office, as shown in Figure 2 a), PDEs and host are coupled and only two states are possible for PDEs: Either the host is turned on and the PDE is FIXED (without resource sharing) or the host is turned off, together with its PDE. In the managed office environment, as shown in Figure 2 b), PDEs are decoupled from hosts and resource sharing is possible. PDEs can perform a transition from FIXED to MOBILE and back, depending on the users behaviour. Additionally, the PDE can be suspended if it is unused (PAUSED) and be resumed when it is needed again. State changes can be initiated by the user or by the system.

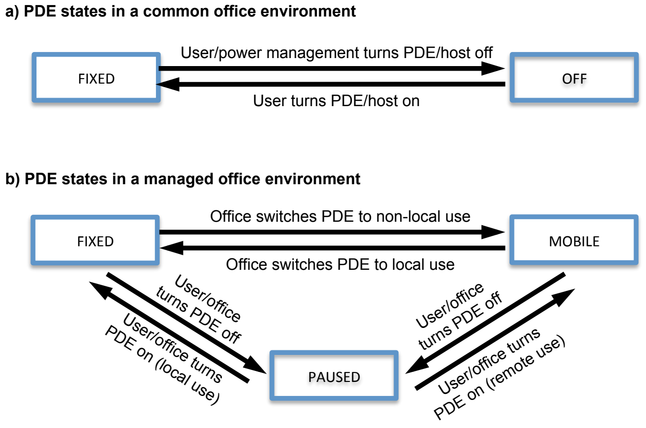

Figure 2: PDEs and their states

Hosts: In the managed office environment each host needs to have the same basic installation. This installation consists of a VMM (Xen) that enables the virtualization of hardware resources. The VMM is extended by a management component that monitors PDEs, manages PDE states, and reports state changes to the ME. Additionally, the host stores a copy of the SPDE, as described in Section 3.2. Hosts can be turned on/off by the ME.

ME: The ME consists of a minimal operating system and an office management service and is responsible for managing hosts and PDEs in the office environment. It monitors their states and manages the overall energy consumption by dynamically suspending/resuming PDEs, redistributing them on hosts and turning hosts on or off. The responsibility area of the ME is configured by the administrator of the office environment and might cover a room, a subnet, or a complete office environment. All hosts are registered with the ME and send update messages, containing the current states of all hosted PDEs.

Management: VMM and ME are cooperating in the management of PDEs and realize an energy-efficient operation based on PDE states. When a PDE is in the state PAUSED, it will be suspended by the VMM, and therefore stops consuming resources. When a PDE is in the state MOBILE, the ME is able to migrate it to another host in order to achieve consolidation and turn off unused hosts. PDE states are either manipulated by users or by VMMs. Users have the possibility to switch between states either by their actions (e.g., logging in remotely) or manually to actively support the energy-efficiency in the managed office. Independent of the manual state switches, the VMM monitors all of its PDEs and changes their states, based on its observations. Within the office, a global variable is set (in the ME) that defines a *critical time period* T_c for two important state changes. First, if a PDE is in the state LOCAL or MOBILE and the virtual CPU of the PDE is idle for T_c minutes, the VMM switches the PDE state to PAUSED. This means the PDEs are unused at the moment and can be suspended. Second, if a PDE is in the state LOCAL and there is no physical user-interaction for T_c minutes, the VMM switches the PDE state to MOBILE. This means, the PDE is used at the moment, but there is no interaction with a local user – the PDE can be consolidated. The critical time period T_c is comparable to the time period that is usually configured for host low-power modes in common offices. Typical values range from 15 to 60 minutes. The effects of T_c on energy efficiency in the common and the managed office are discussed in Section 4. However, similar to the time period of low-power modes, T_c should not be chosen too small, in order to avoid interferences with the work of the user.

The ME determines an energy-efficient mapping of PDEs to hosts in the office and initiates necessary migrations of PDEs. This mapping is called a *configuration* in this paper.

Definition 1: A configuration is called **valid**, if 1) all FIXED PDEs are located at their dedicated hosts, and 2) no host h_i provides more than C_i MOBILE and FIXED PDEs simultaneously.

When a configuration is valid, it is ensured that all users that are working physically at their host, actually have their PDE locally available. Furthermore it is ensured, that none of the hosts h_i in the office provides more than C_i PDEs at the same time, where C_i is the *consolidation factor* of host h_i. C_i determines the maximum number of running PDEs that can be provided by host h_i, based on the available resources of the host. Valid configurations allow all users to access their PDEs as desired, but are not necessarily optimized considering energy efficiency. For simplification it is assumed in this paper that all hosts h_i have a common consolidation factor C.

Definition 2: A configuration is called **host optimal**, when it utilizes the minimum possible number of hosts to provide all required MOBILE and FIXED PDEs in the office.

When a configuration is host optimal, it utilizes a minimum number of hosts h_i to provide all PDEs that are currently used (locally or remotely) in the office. In a host optimal configuration, the number of utilized hosts $H^{on}(t)$ at time t calculates as

$$H^{on}(t) = max \left\{ \mathcal{N}^F(t), \left\lceil \frac{\mathcal{N}^F(t) + \mathcal{N}^M(t)}{C} \right\rceil \right\} \qquad (1)$$

where $\mathcal{N}^F(t)$ is the number of PDEs in the state FIXED and $\mathcal{N}^M(t)$ is the number of PDEs in the state MOBILE in the office.

The ME has to fulfill three goals in the energy efficient mapping of PDEs to hosts.

- The ME needs to constantly maintain a valid configuration in the office environment to provide PDEs to users as needed.
- The ME needs to achieve energy-efficiency through consolidation, by approaching a host optimal configuration.
- The ME needs to minimize the number of migrations within the office environment because migrations are costly themselves (in terms of network traffic and interference with the users work). Unnecessary migrations need to be avoided and hosts should not be overloaded by performing several migrations simultaneously.

In the managed office environment as it is suggested in this paper, the ME restores the validity of the configuration immediately with every state change that occurs in the system. Additionally, to approach a host optimal configuration on the one hand and to reduce the number of migrations on the other hand, the following heuristic is applied:

1. Continue with step 2 as soon as the current configuration utilizes more than $H^{on}(t) + b$ hosts.
2. Select all possible source hosts for PDE migration (hosts that have at least one MOBILE PDE but no FIXED PDE). *Use only hosts that are not already involved in currently performed migrations.*
3. Select all possible target hosts for PDE migration (hosts with free resources). *Use only hosts that are not already involved in currently performed migrations.*
4. Repeat until no further PDE migration is possible:
   (a) Find the source host (in the source-host selection) that contains the smallest number of MOBILE PDEs.
   (b) Find a movable PDE on the source host.
   (c) Find a target host (in the target-host selection). Start search with hosts that contain FIXED PDEs.
   (d) Mark the movable PDE for migration.
   (e) Update source and target hosts.
5. Initiate all of the marked migrations and continue with step 1.

In each iteration, this heuristic approaches a configuration that utilizes $H^{on}(t) + b$ hosts or less, where b is a buffer that calculates as a percentage of the n office hosts. If migrations would last no time (hypothetically), $H^{on}(t) + b$ hosts or less would be utilized. However, migrations last approximately 3-4 minutes [3], therefore the heuristic excludes all hosts that have previously been involved in validation or consolidation processes. This and the buffer b successfully prevent oscillations in the system and reduce the number of migrations. Additionally, to prevent hosts from being overloaded, all migrations that concern the same host are performed strictly sequential. This simple but practical heuristic for the mapping of PDEs to hosts will be enhanced in future work to consider the dynamic of the system in a more detailed way.

4. EVALUATION

The managed office environment as described in Section 3.3 was simulated in a discrete event simulation. The goal of the simulation was to compare the energy consumption of an *unmanaged office (UO)* and a *managed office (MO)* environment, in terms of consumed energy and provided service. The simulation is used to verify two different hypotheses:

1. The suggested energy saving methods in the MO – a) suspending unloaded PDEs and b) consolidating loaded PDEs – are adequate to significantly save energy in an office environment.
2. The service that is provided to the users within the MO is not significantly interrupted. The users are not prevented from doing their day to day work and experience a service comparable to the service in an UO.

UO and MO where simulated over a time period of 12 months, each with 200 users that show similar user behavior in both scenarios. All simulations were initialized by simulating 24 h in advance before taking any measurements. In this section, the energy consumption of office hosts in UOs and MOs is evaluated (following the discussion in Section 2), network costs are not considered.

4.1 Office Environment

Each user has a personal host he usually runs his PDE on. All PDEs are considered to show similar behavior in terms of resource usage. The office shares a common consolidation factor C (see Section 3.1). $C = 1$ represents an office without consolidation, only office-wide energy management is applied, similar to other management solutions as described in Section 5. The optimal value of the buffer b (see Section 3.3) has been determined by simulations and is set to 3% of n. Each host needs 3 minutes to boot/shutdown and each PDE (in the MO scenario) needs 1 minute to be suspended/resumed. The offices are assumed to have Fast Ethernet network, with a throughput of 94 Mbps, which is needed to calculate the transmission time of PDEs. In addition to the plain transmission time, 3 minutes are added to each migration, for synchronizing data between source and target host. The critical time period T_c for the power-management is equal for all hosts and similar in the UO and MO: In the UO the low-power mode of a host is activated after T_c minutes. In the MO PDEs are either suspended or migrated (see Section 3.1) after exactly the same time period. The decision if a host's low-power mode is broken (UO hosts only) is determined following a Bernoulli distribution with parameter $p_{lpbroken}$. The management of the MO is realized by a single (additional) host which is managing the 200 PDEs and hosts. Each host (including the ME) consumes the same amount of energy – 72 Watts if it is "on", 2 Watts if it is "off" (both are average values from Table 1), and 36 Watts in low-power mode (half of the "on" consumption).

4.2 User Behaviour

The user behaviour within the simulated offices was inspired by observations concerning a small office environment (about 20 hosts) at the University of Passau. In future work, other office environments with different user behaviours will be modelled and simulated. It is assumed that users mainly access their hosts within 9 core working hours per day, starting from Monday at 9 h and ending Friday at 18 h. During these hours the users have periods of interaction with their host between 1 minute and 2 hours (exponential distribution, mean 60 minutes), followed by periods of no interaction between 1 minute and 4 hours (exponential distribution mean 30 minutes). This means there are a higher number of short breaks, and a smaller number of longer breaks between periods of host interaction. Each following working day of a certain user starts 15 h after the last one has ended (63 h on Friday afternoon). Outside working hours, users that do not turn off their host (non-energy-efficient users) start overnight jobs (the length of a job varies from 1 h to throughout the night/weekend, following a uniform distribution). After the job has finished, the PDE is idle until the next morning. It is assumed that users that turn their host off manually (energy-efficient users), do so if they leave their host for more than 90 minutes. Users that work remotely (see Section 2), are working remotely for a whole day – remote work and local work does not switch during a single day. The decision if a user works remotely is determined at the beginning of each day following a Bernoulli distribution with parameter p_{remote}. Similarly, the decision if a user turns his host off during the night (energy-efficient user) is determined with parameter p_{off}. To model the costs of migration, as explained in Section 3.2, the size of the PDE user data ($DIFF$) is between 100 MB and 2 GB and follows a normal distribution with mean value 1 GB and a variation of 500 MB. Migrations cause overhead in terms of running hosts in the evaluation – when a PDE is transferred from host to host, both of the hosts have to remain powered on during the migration.

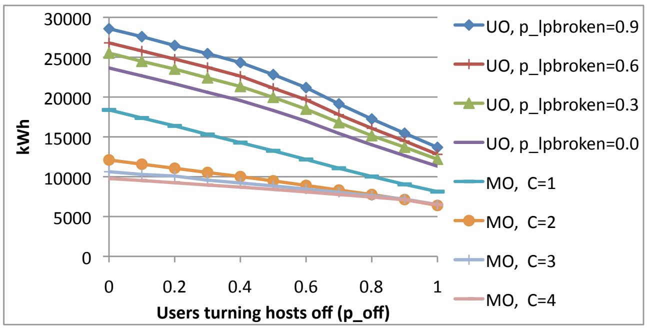

Figure 3: Energy consumption of an office environment with $p_{remote} = 0.25$

4.3 Results

In Figure 3 the energy consumption of the UO is compared to the energy consumption of the MO, where on average 25% of all users worked remotely. The x-axis shows the mean ratio of energy-efficient users that turn off their host over nights. The y-axis shows the consumed energy over 12 weeks. The top 4 curves illustrate the energy use of an UO with different percentages of non-functional low-power modes (see Section 2). As expected, an UO with a mean of 90% broken low-power modes consumes the most energy in this simulation. It can be observed that the energy consumption of the UO decreases, the more users show energy-efficient behavior. The lower 4 curves show the energy consumption of a MO considering different consolidation factors. $C = 1$ means in this case that no consolidation of PDEs takes place at all – only an office-wide energy management is applied, comparable to power-management solutions for environments, as described in Section 5). The curve illustrates the savings of energy that can be achieved, just by automatically turning off unused PDEs (together with their hosts) after a critical time T_c (which is 45 minutes in this case). It can be observed that by applying this kind of management a significant reduction of energy is achieved. When the consolidation factor is raised to 2, additionally consolidation is done in the system and PDEs that are not used locally are migrated to other hosts. Again, a significant portion of energy is saved by this method in the MO, even a higher portion of energy is saved in this step, than by pure management ($C = 1$). It can be observed, that a further increase of the consolidation factor ($C = 3$ or $C = 4$) does not lead to further significant energy savings in this scenario. There are two main reasons for this effect. 1)

In the shown scenario 75% of all users (local users) need a dedicated host during the working day. The remote users can easily be distributed to these hosts with $C = 2$ and a higher consolidation factor doesn't change the situation. 2) Higher consolidation factors impose more overhead to the system (in terms of running hosts), thus decreasing the benefit. This is due to the simple consolidation heuristic as it is described in Section 3.3. As described, migrations are performed strictly sequentially if they concern the same host. If several PDEs need to be migrated to (or from) a single host, PDEs have to wait for migration. During this waiting time, both hosts are "on" (the source and the target host). This effect raises with a higher consolidation factor, because the consolidation concerns fewer hosts.

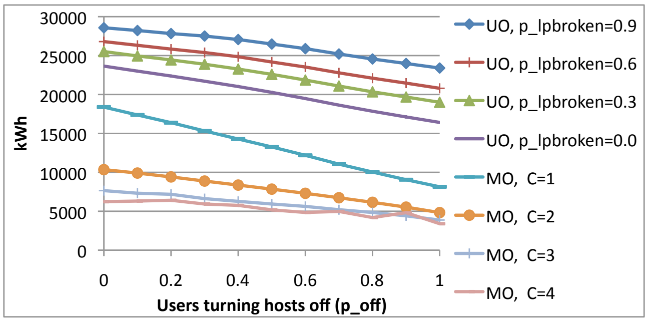

Figure 4: Energy consumption of an office environment with $p_{remote} = 0.75$

Energy savings increase, when the proportion of remote workers increases, this can be observed in Figure 4. It shows basically the same graphs than in Figure 3, however, with a mean of 75% of remote workers in the office environment. The most interesting fact in Figure 4 is that a high number of remote users increases the savings that are created by consolidation ($C = 2$ and $C = 3$), whereas the savings achieved by management stay nearly the same ($C = 1$). A high number of remote users leads to significant savings in the MO, whereas the UO shows a trend of using rather more energy – in the UO remote users leave their host on over night to be able to use it again remotely in the next morning.

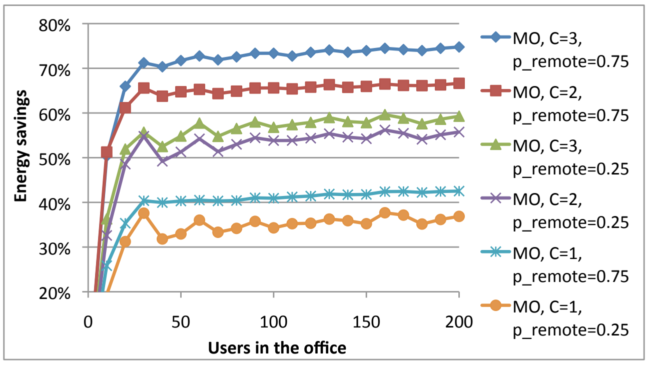

Figure 5: Energy savings with an increasing number of users

Energy saving are also illustrated in Figure 5. In this simulation all parameters are set as in the previous simulation. The average percentage of users that turn their host off (energy-efficient users) is fixed to 40% and the average percentage of non functional low-power modes is fixed to 60%. The x-axis shows a growing number of users in the office and the y-axis shows the energy savings that were achieved by the MO approach. It can be observed, that the highest number of remote users leads to the highest energy savings when consolidation is enabled ($C = 2$ and $C = 3$). The suggested MO saves up to 75% of energy in the illustrated scenario. Without consolidation ($C = 1$) only up to 43% of energy savings are achieved. Even offices with a small number of employees can significantly save energy, depending on the user behaviour.

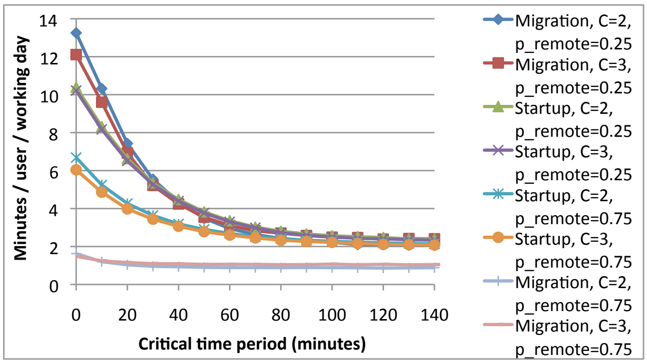

Figure 6: Migration times and startup times (per user and per working day)

Figure 6) compares the UO and the MO in terms of the service which is provided to the users in the office. In the simulation all parameters are set similar to the previous simulation. The number of users is fixed to 200. The x-axis denotes the critical time period T_c after that a PDE is suspended or migrated (see Section 3.3). The y-axis illustrates waiting times of users in the MO. The migration curves illustrate the times that users have to wait until they are able to access their PDE locally. Effects on the user during the migration time depend on the type of migration (see Section 3.2) that is implemented in the system. An implementation of the described transient thin-client mode would, for instance, lead to a slightly decreased quality of service during this time. The startup curves illustrate waiting times according to the startup of PDEs and hosts and are more important to the user. During these times, the user has actually to wait for the service to be provided. It can be observed that the waiting times decrease with an increasing number of T_c, due to unnecessary suspension or migration of PDEs. With a T_c of about 45 minutes, the startup waiting times per user and day reduces to less than 4 minutes for $p_{remote} = 0.25$ and to less than 3 minutes for $p_{remote} = 0.75$ per day and user in this scenario, which seems to an acceptable trade-off, considering the achieved energy savings.

5. RELATED WORK

The development of energy efficient IT equipment is fostered by labels such as the US Energy Star[8] or the European TCO Certification[9] which rate IT products (mostly monitors) according to their environmental impact. Novel emerg-

---

[8] http://www.eu-energystar.org

[9] http://www.tcodevelopment.com

ing technologies such as solid-state disks consume much less energy than the currently used hard-disk drives. Computer power can be saved by means of various well-known techniques. First, the processor can be powered down by mechanisms like SpeedStep [11], PowerNow, Cool'n'Quiet or Demand-Based Switching. These measures enable slowing down CPU clock speeds (clock gating), or powering off parts of the chips (power gating), if idle [12]. By sensing lack of user-machine interaction, different redundant hardware parts can incrementally be turned off or put in hibernating mode (display, disk, etc.). The Advanced Configuration and Power Interface (ACPI) specification [9] defines four different power states that an ACPI-compliant computer system can be in. All of these techniques attempt to minimize the power consumption of a single device, managed individually by a user. In contrast, this paper focuses on the office environment as a whole, exploiting centralized power management policies and a globally managed consolidation of resources.

There are several projects that provide power-management solutions for office environments. Examples are eiPowerSaver[10], Adaptiva Companion[11], FaronicsCore[12], KBOX[13], or LANrev[14]. In such approaches, office-wide power management policies are applied to office environments. Office hosts change to low-power modes, independent of user-specific power management configurations. Additionally, mechanisms are provided to wake up hosts if necessary. This way, hibernated hosts can be used for overnight jobs (e.g., backup processes) and for remote usage. Such solutions, however, still rely on the capability of the host to switch to low-power modes. This capability depends on the complex interaction of a host's hard and software. The approach presented in this paper is independent of such interaction. PDEs are suspended together with their VM without being aware of it. Office hosts, on the other hand, have a common simple software configuration, as explained in Section 3.3 which makes a low-power configuration easier. What is more, the mentioned power-management solutions focus on idle hosts only. The solution suggested in this paper, additionally deals with the energy consumption of underutilized hosts in office environments. A comparison of both approaches in terms of saved energy is shown in Section 4.

The term cloud computing [8] has been introduced recently and refers to data-centre-based services, stored in ubiquitous computing clouds and is strongly related to grid computing [17]. Cloud computing approaches try to offer computing power independent of actual hardware location. In a cloud, scalable and virtualized hardware resources are provided as a service. VMs are running in a distributed environment and can be migrated to hardware that currently provides idle resources. Popular clouds are, e.g., Amazons Elastic Compute Cloud (EC2), or Google App Engine. In particular, Cloud Computing is an inherently energy-efficient virtualization technique, in which services run remotely in a ubiquitous computing cloud that provides scalable and virtualized resources. Thus peak loads can be moved to other parts of the cloud and the aggregation of a cloud's resources can provide higher hardware utilization. In contrast to the approach presented in this paper, virtualization and consolidation in clouds focus on highly centralized and controllable high-performance data-centre environments. Such environments usually consist of homogeneous hardware which is located close to each other in racks, interconnected via a high performance networks, and administrated by a small group of persons. This paper, on the other hand, focuses on the energy efficiency of office environments – outside of the data centre – where a high number of heterogeneous hosts are typically connected via Fast Ethernet and are directly accessed by a high number of users.

Thin-client/terminal-server approaches uses data-centre technology to provide energy-efficient services in office environments. User environments (operating systems, applications, and data) are provided by terminal servers and users can access these environments via energy-efficient thin clients. Common terminal-server software products are Citrix XenApp[15], Microsoft Windows Server 2008[16], or the Linux Terminal Server Project[17]. Similar to the approach suggested in this paper, such approaches foster a resource sharing among users in the office environment. A comparison of both approaches in terms of energy consumption will be done in future work. A clear difference is, however, that the thin-client/terminal-server approach is based on the installation of additional devices within and outside of the data centre (energy-efficient thin clients and terminal servers). Instead, the approach suggested in this paper utilizes available hosts in office environments to enable resource sharing.

In [10, 4] a virtualized future home environment is introduced that uses virtualization to aggregate and consolidate distributed hardware resources of home users in order to save energy. Similar to offices, also in home environments some machines are running on a 24/7 basis (e.g., media servers or P2P clients). These services can be consolidated by using different virtualization techniques in order to turn unused hosts off. In contrast to the future home environment approach, this work focuses on resource sharing in office environments as they can be found today in companies or public administration. Whereas in the future home environment separate services are virtualized (e.g., video-encoding or P2P file-sharing services) and are distributed among homes, this work suggests to virtualize user environments (PDEs) as a whole. As an important consequence, the approach in this paper envisions a seamless and transparent provision of user services within the PDE (e.g., when a PDE has been migrated, the user still finds his text document open, with the cursor at the same position as before the migration). The future home environment approach, in contrast, is not transparent to the user. The user has to utilize special software that enables the envisioned migrations of services, and seamless access to migrated services is not possible. Instead the result of a service is transferred back to the user.

6. CONCLUSIONS AND FUTURE WORK

This paper has analyzed the potential of energy savings in office environments and discussed methods and requirements of resource management in offices to exploit this po-

---

[10] http://entisp.com/pages/eiPowerSaver.php

[11] http://www.adaptiva.com/products_companion.html

[12] http://faronics.com/html/CoreConsole.asp

[13] http://www.kace.com/solutions/power-management.php

[14] http://www.lanrev.com/solutions/power-management.html

[15] http://www.citrix.com/XenApp

[16] http://www.microsoft.com/windowsserver2008

[17] http://www.ltsp.org

tential. An architecture based on virtualization methods has been presented that shares and manages resources in office environments. The proposed solution extends available power-management approaches and is opposed to datacentre based thin-client/terminal-server solutions. A shift from current decentralized resource management approaches (per user) is suggested to a centralized resource management approach (per office). Simulation results indicate that 75% of energy savings can be achieved by the suggested approach. It has also been shown in the paper that such savings can be achieved without significantly decreasing the quality or quantity of services in the office environment.

In future work, the suggested architecture has to be refined and different user scenarios have to be analyzed. Also resilience and security issues have to be considered.

7. ACKNOWLEDGEMENT

The research leading to these results has received funding from the German Federal Government BMBF in the context of the G-Lab_Ener-G project and from the European Community's FP7 in the context of the EuroNF Network of Excellence (grant agreement no. 216366).

8. REFERENCES

[1] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield. Xen and the art of virtualization. *SIGOPS Oper. Syst. Rev.*, 37(5):164–177, 2003.

[2] F. Bellard. QEMU, a fast and portable dynamic translator. In *Proceedings of the USENIX Annual Technical Conference, FREENIX Track*, pages 41–46, 2005.

[3] A. Berl and H. de Meer. An Energy-Efficient Distributed Office Environment. In *Proceedings of European Conference on Universal Multiservice Networks (ECUMN 2009), Sliema, Malta, October 11-16, 2009*. IEEE Press, October 2009.

[4] A. Berl, H. Hlavacs, H. de Meer, and T. Treutner. Virtualization Methods in Future Home Environments. *IEEE Communications Magazine*, December 2009.

[5] P. Bertoldi and B. Atanasiu. Electricity consumption and efficiency trends in the enlarged European Union. *IES–JRC. European Union*, 2007.

[6] C. Cartledge. Sheffield ICT Footprint Commentary. *Report for SusteIT. Available at: http://www.susteit.org.uk/files/files/26-Sheffield _ICT_Footprint_Commentary_Final_8.doc*, 2008.

[7] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield. Live Migration of Virtual Machines. In *2nd conference on Symposium on Networked Systems Design & Implementation (NSDI'05)*, pages 273–286, Berkeley, CA, USA, 2005. USENIX Association.

[8] C. Hewitt. Orgs for scalable, robust, privacy-friendly client cloud computing. *IEEE Internet Computing*, 12(5):96–99, 2008.

[9] Hewlett-Packard. Microsoft, Phoenix, and Toshiba. Advanced configuration and power interface specification. *ACPI Specification Document*, 3, 2004.

[10] H. Hlavacs, K. A. Hummel, R. Weidlich, A. Houyou, A. Berl, and H. de Meer. Distributed Energy Efficiency in Future Home Environments. *Annals of Telecommunication: Next Generation Network and Service Management*, 63(9):473–485, October 2008.

[11] Intel. White paper 30057701. wireless intel speedstep power manager: Optimizing power consumption for the intel pxa27x processor family, 2004.

[12] Intel and U.S. Environmental Protection Agency. Energy star* system implementation, whitepaper. *SIGCOMM Comput. Commun. Rev.*, February 2007.

[13] J. Koomey. Estimating total power consumption by servers in the US and the world, Technical report. Technical report, Lawrence Berkeley National Laboratory Stanford University, February 2007.

[14] M. Satyanarayanan, B. Gilbert, M. Toups, N. Tolia, D. O'Hallaron, A. Surie, A. Wolbach, J. Harkes, A. Perrig, D. Farber, et al. Pervasive personal computing in an internet suspend/resume system. *IEEE Internet Computing*, pages 16–25, 2007.

[15] R. Steinmetz and K. Wehrle. *Peer-to-Peer Systems and Applications (LNCS)*. Springer-Verlag New York, Secaucus, NJ, USA, 2005.

[16] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *SIGCOMM '01*, pages 149–160. ACM Press, 2001.

[17] L. Vaquero, L. Rodero-Merino, J. Caceres, and M. Lindner. A break in the clouds: towards a cloud definition. *ACM SIGCOMM Computer Communication Review*, 39(1):50–55, 2008.

[18] C. Webber, J. Roberson, M. McWhinney, R. Brown, M. Pinckard, and J. Busch. After-Hours Power Status of Office Equipment in the USA. *Energy-the International Journal*, 31(14):2487–2502, 2006.