Network Virtualization in Energy-Efficient Office Environments

Andreas Berl[a], Nicholas Race[b], Johnathan Ishmael[b], Hermann de Meer[a]

[a] *University of Passau, Innstr. 43, Passau, Germany*
[b] *Lancaster University, InfoLab21, Lancaster, UK*

Abstract

The rising costs of energy and world-wide desire to reduce CO_2 emissions has led to an increased concern over the energy efficiency of information and communication technology. Whilst much of this concern has focused on data centres, office environments (and the computing equipment that they contain) have also been identified as a significant consumer of energy. Office environments offer great potential for energy savings, given that computing equipment often remains powered for 24 hours per day, and for a large part of this period is underutilised or idle. This paper proposes an energy-efficient office management approach based on resource virtualization, power management, and resource sharing. Evaluations indicate that about 75% energy savings are achievable in office environments without a significant interruption of provided services. A core element of this office management is a peer-to-peer network that interconnects office hosts, achieves addressing and mediation, and manages energy efficiency within the office environment. Several peer-to-peer approaches are suggested and discussed in this paper. Two of the approaches are evaluated, based on a discrete event simulation.

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

1. Introduction

Information and communication technology costs have seen a shift in recent years, with the rapid reduction in hardware costs contrasted with the steep increase in energy costs. Koomey [1], reports that data centers in the USA and worldwide have doubled their energy consumption from 2000 to 2005. In addition,

Email addresses: berl@uni-passau.de (Andreas Berl), n.race@lancaster.ac.uk (Nicholas Race), j.ishmael@lancaster.ac.uk (Johnathan Ishmael), demeer@uni-passau.de (Hermann de Meer)

end-devices have also considerably contributed to the increase of power consumption, according to a 2006 survey [2] commissioned by the EU.

Hosts in office environments are often running without being locally used, i.e., the host is not physically accessed by a user. This happens for short time periods (e.g, if users are in meetings, make telephone calls, have lunch or coffee breaks, etc.) as well as for longer periods of time. Physically unused hosts are often left switched on, because users require access to them remotely. Remote access typically happens from the user's home or when users are working externally. Remote access is needed in such cases to access applications and data in the office. The user may need access to email accounts, personal data, or applications. Another important cause that leads to physically unused but running hosts are overnight jobs. A user might schedule a job (e.g., a simulation, a download or a backup) outside his working hours, so it 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. [3] have analyzed sixteen sites in the USA and reported that 64% of all investigated office hosts were running overnight.

Locally unused hosts provide a huge potential for energy savings. Either such hosts are idle (0% CPU load), used remotely by their user, or performing a job without user interaction. In all cases energy can be saved. Idle hosts consume a considerable amount of energy, compared to computers that are turned off, without providing any added benefit. Measurements that have been performed at the University of Sheffield on hosts that are typically used as personal computers [4] show that idle hosts still consume 49% to 78% of the energy that they need when they are intensely used. Such hosts have to be stopped from consuming resources. If a locally unused host is not idle (i.e., is used remotely or performs a job) there is also a potential to save energy. Such hosts are often underutilized by typical office applications (e.g., text processors, browsers, or mail clients), leading to a high number of lightly utilized hosts that consume nearly as much energy as the same number of heavily utilized hosts. However, only users with physical host access need a separate host to work with, other users do not necessarily need to utilize separate hosts. Local users should share their resources with non-local users to increase the utilization of hosts.

Several approaches have been suggested that deal with high energy consumptions of hosts in office environments (see Section 6). 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 distributed 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 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 describes a managed office environment based on virtualization that achieves energy efficiency in offices. Section 3 presents virtualization approaches to enable a managed office environment and Section 4 discusses three peer-to-peer overlay approaches in detail. Section 5 evaluates energy consumption and overhead of the suggested peer-to-peer approaches within a common office environment (based on the office environment of the University of Sheffield [4]). Section 6 discusses related approaches and Section 7 concludes this paper.

2. A Managed Office Environment

When a user powers on his host in a common office, he finds his usual working environment: within this paper we refer to this working environment as a *personal desktop environment (PDE)*. This typically consists of an operating system, applications, and the user's personal configuration. Although roaming profiles are often available in common offices (see Section 3), 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, the host is also turned on/off and vice versa. Users are able 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 (between physical hosts) 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 remotely utilizes his PDE – the PDE may be provided by a different host. Mobile services can be achieved by using virtualization methods, as explained in Section 3.

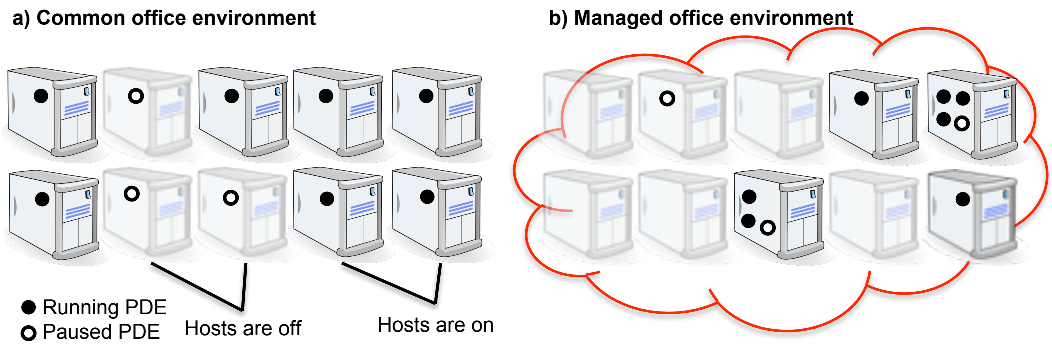

Figure 1: Common and managed office environment

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. As can be seen in the figure, 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:

- Unloaded PDEs in the office environment are halted, thus stopping them from consuming resources. If a PDE is idle (no job is performed on behalf of its user) it will be suspended.
- Loaded 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.

These steps provide the possibility to optimize the energy consumption of each single host $\sum_{i=1}^{n} min(E(h_i))$, where n is the number of hosts $(1 \leq i \leq n)$ and $E(h_i)$ is the energy consumption of host h_i. Additionally, it is possible to reduce of the overall energy consumption $min(\sum_{i=1}^{n} E(h_i))$ within the office environment, by considering the office as a whole.

The managed office environment achieves an energy-efficient provisioning of office services by utilizing available office hardware (office hosts and network). In contrast to other approaches (e.g., cloud computing or terminal server approaches, see Section 6) there is no need to acquire additional hardware as thin clients or data centre equipment and, moreover, all services remain inside of the office. The energy management needs to achieve energy efficiency within the office while being mostly transparent to the users. To provide a service similar to common offices, each PDE needs to be moved to its user's office host when the user powers it on. On the other hand, any host that provides enough idle resources can serve remotely working users. The managed office environment has to dynamically determine an energy-efficient mapping of PDEs to hosts and to initiate necessary migrations of PDEs. This mapping is further described in [5]. It has to fulfill contradicting goals and needs to solve a multidimensional optimization problem:

- The mapping needs to constantly maintain a **valid** configuration in the office environment to provide PDEs to users as needed. A mapping is called valid, if 1) all PDEs are located at their dedicated hosts, and 2) no host is overloaded with PDEs. Valid mappings allow all users to access their PDEs as desired, but are not necessarily optimized considering energy efficiency.
- The mapping needs to achieve energy efficiency through consolidation, by approaching a **host optimal** configuration. A mapping is called host optimal, when it utilizes the minimum possible number of hosts to provide all required PDEs (locally or remotely) in the office.
- The mapping 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 user's work). Unnecessary migrations need to be avoided and hosts should not be overloaded by performing several migrations simultaneously.

3. Virtualization of Office Resources

3.1. Requirements

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 PDEs 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 sufficintly flexible 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 unused hosts to power off (to save energy) and then powered on again if they are needed 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 hosts within the office environment have to be logically connected in order to enable a mediation of free resources and PDEs. In traditional 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 tolerated by users.

3.2. System virtualization

An important first step in the approach to virtualizing a managed office environment is system virtualization. It enables service consolidation and is successfully deployed in data centres today. It can be adapted to office environments in order to achieve a similar utilization and energy savings of office resources. In system virtualization *virtual machines (VMs)* are created from idle resources. Full hosts are virtualized, consisting of virtual CPUs, memory, hard disks, network interface cards, 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 a *Virtual Machine Monitor (VMM)* (e.g., VMWare Server[1], QEMU [6], or Xen [7])) and is able to process several VMs simultaneously on a single host. System emulators (e.g., QEMU) emulate their own hardware environment and translate dynamically between different Instruction Set Architectures (ISA) such as x86 or SPARC. Without emulation, VMs are limited to the ISA provided by the physical hardware, while emulation enables any combination of host and guest ISA. However, the translation process causes considerable overhead. To support heterogeneous office equipments (in the managed office environment approach) without imposing emulation overhead, it is possible to create separate PDEs for each ISA. The PDE management needs to consider the different architectures in its PDE-host mapping and to consolidate PDEs only on the appropriate machines. 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 *live migration* [8]. 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 2, can be encapsulated within a VM and inherits all of the VM-related features. Therefore, system virtualization, together with the management functions concerning VMs meets several of the requirements stated in Section 3.1. It enables the operation of PDEs in separated run-time 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[2], 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 6) is that it works completely independently 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 an interruption of running services.

However, the costs of migration (as discussed in [9]) represent 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 in their entirity. This leads to considerable overhead because operating system, user data and applications represent 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

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

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

preconfigured full featured operating system (e.g., Windows or Linux), together with common applications, that provides the basis for each PDE. 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, the complete PDE is not transferred. Instead, the difference $DIFF = PDE - SPDE$ (e.g., calculated by the rsync[3] tool) 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 [10]. 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.

Apart from the differences in terms of host and network performance between office environments and data centres, the resource management is also more complex. Whereas the data centre is a controlled environment where only administrators have physical access to hosts, the office environment is rather uncontrolled. Users are able to power hosts on and off, unplug cables, or move hosts to other locations. Furthermore, in data centres users access their services remotely, which eases up the consolidation of services. Local access to hosts, as is typical in office environments adds hard constraints to the management of resources: Services that are used locally cannot be migrated or consolidated.

[3] http://rsync.samba.org/
[4] http://www.samba.org/
[5] http://www.openafs.org/

3.3. Peer to Peer overlays

A second important virtualization approach that is required in order to realise the managed office environment is the use of peer-to-peer (P2P) technology. Independent of the logical network that is used to interconnect hosts, the resource sharing in the managed office environment is achieved in a 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 and principles from P2P overlays can be used to realize a management environment 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, the behaviour of which can be adapted to office environments. First, these kinds of network 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). The following section suggests different P2P-based approaches to realize the suggested managed office environment.

4. P2P-based Management Environment

The suggested office environment is managed by a *management environment (ME)*. The ME monitors states of PDEs and hosts in the office and manages the overall energy consumption by dynamically suspending/resuming PDEs, redistributing them between 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

[6] http://www.overnet.org/
[7] http://www.bittorrent.com

the current states of all hosted PDEs. The ME determines an energy-efficient mapping of PDEs to hosts (see Section 2) in the office and initiates necessary migrations of PDEs. Three different approaches of P2P overlays are possible to realize the suggested ME in the office environment: 1) centralized client/server-based 2) pure P2P-based, and 3) a hybrid approach.

4.1. Centralized P2P-overlay approach

The most simple approach in terms of setup, administration, and management is the centralized tracker based approach [11]. In this configuration, one dedicated server provides the necessary management channel through which state changes can be issued, and feedback communicated. All hosts within the office environment are logically connected to the centralized server, as illustrated in Figure 2. Although the ME is client/server-based, resource-sharing is still performed by office hosts in a distributed P2P manner. The peers of this P2P overlay are the hosts of the office environment, the PDEs themselves are not aware of the P2P network.

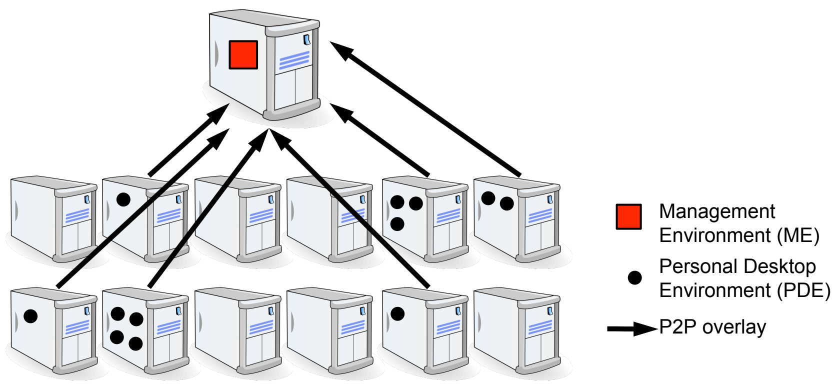

Figure 2: A centralized P2P-overlay approach

The use of a centralised ME, mimicking a tracker from traditional P2P file sharing networks has several benefits. The system is non-complex to design and implement, as it is not necessary to deal with state replication amongst different management servers. This also means there is only one server to administer the PDEs and their hosts. As the ME is a constant it can be allocated a static IP address, reachable from outside the office environment, allowing remote-workers access to their PDEs. Also, the ME is located on a dedicated hardware platform, therefore, it is likely to be kept in a secure location. This reduces the risk of failure due to accidental disruptions (e.g., accidentally removing the power to a host).

Similarly this centralised/tracker based approach has a number of disadvantages. The main one of which is the fact that additional (energy consuming) hardware is needed to realize the management platform, somewhat in contrast to the purpose of this scientific work. As the number of PDEs on-line reduces (e.g., at the end of a working day) the energy resources associated with managing them do not reduce. Having the management functionality in a single place introduces all commonly known client/server architecture problems – such as the risk of a single point of failure or issues with scalability, should there be a large number of office machines in use. Any failure of the ME would result in failures of PDEs. Sub-environments, say a department within a building would be difficult to create: it may require a new ME server (together with additional energy resources).

While a dedicated ME server is a simplistic approach, it is in contradiction to the energy-efficient approach. The ME server would have to be a dedicated host with supporting hardware (networking equipment, UPS). The power this machine consumes would waste additional energy within the office environment, contrary to the goal we have set out to achieve.

4.2. Pure P2P-overlay approach

A pure P2P-overlay approach (e.g., Chord-based [12]) doesn't rely on a central server (a tracker in traditional P2P terms) to provide management of the overlay. Instead, it relies upon information sharing amongst nodes to build up an accurate picture of the situation. Applying a pure P2P-overlay approach to the distributed office environment is a complex task, in which hosts are responsible for managing themselves. Hosts need to provide PDEs and a management instance, in this case. As no global view of the office environment is available, the global goal of energy-efficiency has to be achieved by distributed management algorithms. An example (based on Chord) of a pure structured P2P overlay is illustrated in Figure 3. Not all of the necessary overlay links are depicted in the Figure.

The main advantage of a pure P2P approach is that of scalability and robustness in the light of node failure. As there is no centralised ME, the system is theoretically able to scale to any size of office environment. Each ME takes a small volume of the management load and provides a localized view on the overlay to achieve an approximation of a global view in the system.

Unlike in a centralised approach, an ME failure will have limited impact (PDEs will have to associate with a new management entity). As MEs would only be embedded on running hosts, they would only

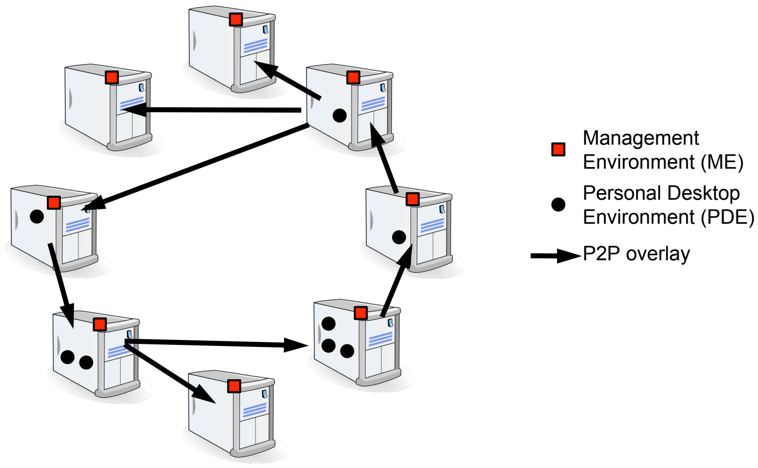

Figure 3: A pure P2P-overlay approach

consume a small amount of energy. Additionally, as the number of running PDEs is reduced, also MEs are shut-down. Therefore, the overlay management requires less resources and scales with the number of running PDEs.

The potential benefits come at the cost of significant implementation complexity [13] and possible inaccuracies with the state of the overlay. Each ME must maintain information on the location of other MEs, PDEs and their state. As there is no central body to manage decisions typically a voting process must be used for every decision (e.g., shut-down or start-up of a physical host). Sub-environments within an environment are also complicated to realize, as there is no central authority it would be difficult to logically divide the network. Another disadvantage to this approach relates to security, as MEs are co-located with PDEs of users and run on the same hosts, users potentially have direct access to ME virtual machines. Finally the fully distributed management is likely to have an overhead and is probably less accurate, due to a missing global view and longer message propagation times [14]. Also typical bootstrapping problems have to be solved, along with the problem of connecting from outside the office (where a central node with a fixed IP is missing).

While the pure P2P approach provides a number of advantages over the centralised approach including robustness and resilience, the pure structured approach is the most complex approach and therefore most difficult to implement.

4.3. Hybrid P2P overlay approach

A hybrid P2P overlay approach is a compromise between the centralized and the pure P2P approach. This design simulates a super-peer [15] approach within P2P networking, in which several clients with suitable resources (CPU/Memory) take on a management role within the P2P network. Within the managed office environment one or more management entities are enclosed within VMs, similar to PDEs and are responsible for managing a local segment of the office environment (e.g., a department). The MEs are available to migrate between different physical hosts and exchange state information amongst other MEs. The use of the P2P hybrid approach enables a managed office environment without imposing a need for further infrastructure to support it; however, the physical hosts will have the overhead of a number of additional VMs (MEs) to support on the physical hardware.

The hybrid P2P overlay is illustrated in Figure 4. The MEs provide management of the P2P overlay and are co-located, but are separate entities on physical hosts with PDEs. Unlike the pure P2P approach, the PDEs are not aware of the P2P network, reducing the complexity of the system. All MEs are interconnected with each other (organization layer overlay) providing state replication and management amongst MEs, as well as permitting fail over should a host running a ME fail.

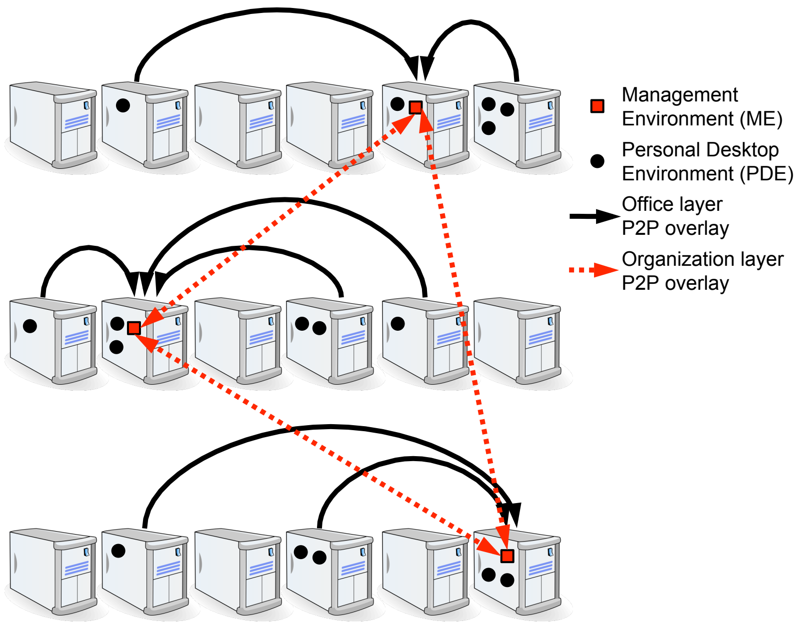

Figure 4: A hybrid P2P-overlay approach

The main advantage of this approach is that of energy efficiency, as the ME service can be consolidated on existing hardware. This approach also provides redundancy in the form of multiple MEs, yet has a relatively low complexity and so is easy to implement and manage. Similar to PDEs, the management entities are movable and can therefore be consolidated with other PDEs in order to save energy.

A hybrid approach also has a number of other advantages including the ability to separate environments into logical sub-environments, each ME maybe responsible for different departments, with different legal requirements and so with different policies on machine roaming. Such an approach would also permit easy administration and management of each sub-environment. The approach is also scalable to any size of office environments as MEs can be added to service additional PDEs as required. This approach also provides robustness, as MEs maintain a distributed state. If an ME fails, other MEs are able to resume the current status of the sub-environment and to establish a new ME. Finally the MEs in the VM consume only a small portion of energy, compared to the server approach. Reliability of hosts and the information they store has always been a significant issue within P2P networks, especially in fully decentralized networks. The common method of overcoming unreliable hosts is to adopt a super-peer structure. Nodes that are identified as stable (using metrics including high system up-time) act as coordination points and data repositories. The super-peers maintain a secondary overlay above the network maintaining state information of the overlay network and its nodes. The hybrid approach, as it is suggested in this section, has similarities to a super-peer approach. In order to improve reliability, selected nodes (managed by the MEs) may act as replication nodes within the network. Should an underlying host fail, any recent modifications to a PDE are available from one or more hosts within the network. Alternatively, if a PDE is moved between hosts, the original copy of the PDE could be maintained allowing minimal version changes to be passed back, keeping the copy up-to-date. Should a host holding a copy go off-line to save power, the ME could hold version changes until the host is back on-line. In an unmanaged office, when a host breaks down, the user may lose his PDE and all the data contained. Within a managed office there is a high probability that there are other (possibly outdated) copies of the PDE available.

The hybrid approach introduces a number of challenges relating to the network management. Firstly the MEs must manage their own P2P network, determining where they should be located and what hosts they should be responsible for. They also have to consider the PDEs and their distribution and management. Additionally, like the pure P2P approach as MEs are co-located with PDEs and run on the same hosts, users have the potential to disrupt the network (by turning off a machine, for example). Similarly MEs may be impacted by performance issues, depending on the load of PDEs that run on the same host. Fail-safe mechanisms must also be developed, to deal with the situation when an ME fails, which ME is responsible for taking ownership of any abandoned PDEs and who is responsible for creating a new ME. Finally, as the MEs are mobile, there may be difficulty in finding an ME when bootstrapping into the network for the first time (or connecting from outside the office environment). A commercial service such as DynamicDNS[8] or academic techniques such as those provided by Atkinson, et al. [16], in which a user updates their DNS record whenever they move, may be suitable for this purpose.

The use of a hybrid approach in a distributed office environment provides a suitable balance between achieving as much power saving as possible, providing greater flexibility whilst reducing implementation complexity when compared to the pure P2P approach.

5. Evaluation

The managed office environment (as described in Section 2) was simulated in a discrete event simulation. The goal of the simulation was to compare the energy consumption of a *unmanaged office (UO)* and a *managed office (MO)* environment, in terms of consumed energy and provided service. The MO has been implemented in two ways, in a centralized client/server based approach (MO_c) and in a hybrid approach (MO_h), both approaches are described in Section 4. The simulation verifies the following 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 MO_h approach provides an energy efficient office management similar to the MO_c approach, while consuming less energy for management efforts.
3. Users in the MO 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 a similar user behaviour in the office scenarios. All simulations were initialized by simulating 24 h in advance

[8] https://www.dyndns.com/

before taking any measurements. The consolidation algorithm of the MO was implemented as a simplified heuristic in these simulations and and will be extended to a more comprehensive algorithm in future work. Only the energy consumption within the office environments is considered in the simulations. The energy consumption outside of the office environment (caused by remote access) remains identical for UO and MO, because the simulations are based on a similar user behaviour.

The following assumptions were made concerning the office environments: Each user has a personal host he usually runs his PDE on. All PDEs are considered to show similar behaviour in terms of resource usage. The office shares a common consolidation factor C, where each host is able to run C PDEs simultaneously (see Section 3). $C = 1$ represents an office without consolidation, only office-wide energy management is applied, similar to other management solutions (see Section 6). Each host needs 3 minutes to boot/shutdown and each PDE (in the MO scenarios) needs 1 minute to be suspended/resumed. The offices are assumed to have a Fast Ethernet network, with a throughput of 94 Mbps, which is used 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. Most current hosts provide low-power modes that can be configured by the user and kick in when a host is idle. The *critical time period* T_c that triggers this power management in the UO and MO is set to 45 minutes. If a PDE is idle for T_c minutes in the MO, it will be suspended and if it is locally unused for T_c minutes, it becomes movable and can be moved to other hosts for energy reasons. Typical values T_c (in common offices) range from 15 to 60 minutes. T_c should not be chosen too small, in order to avoid interference with the work of the user. Many devices that are low-power capable do not successfully enter such modes (e.g., only 4% of functional low-power modes are reported in [3]). Low-power modes are subject to the complex combined effects of hardware, operating systems, drivers, applications and the user-based power management configuration. The decision if a UO host's low-power mode is broken is determined following a Bernoulli distribution with parameter p_{lpb}. The energy consumption of the office environment is chosen according to the office environment of the University of Sheffield and is based on data given in the report Sheffield_ICT_Footprint [4]. The University of Sheffield has examined personal computer purchase records over a period of 5 years from three main suppliers, which showed over 10,000 purchases, mostly stationary hosts, portable hosts represented only a small portion of these purchases. Based on this data, each host (also the centralized server ME) consumes the same amount of energy – 72 Watts if it is “on”, 2 Watts if it is “off”, and 36 Watts in low-power mode (half of the “on” consumption).

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 without 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. Users that work remotely, 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_{rem}. Similarly, the decision if a user turns his host off if it is unused (energy-efficient user) is determined with parameter p_{off}. It is assumed that energy-efficient users turn their host off manually if they leave their host for more than 90 minutes. To model the costs of migration, as explained in Section 3, 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.

In Figure 5 the energy consumption of an UO is compared to the energy consumption of the MO in terms of KWh used, where on average 25% (a) or 75% (b) 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. 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,

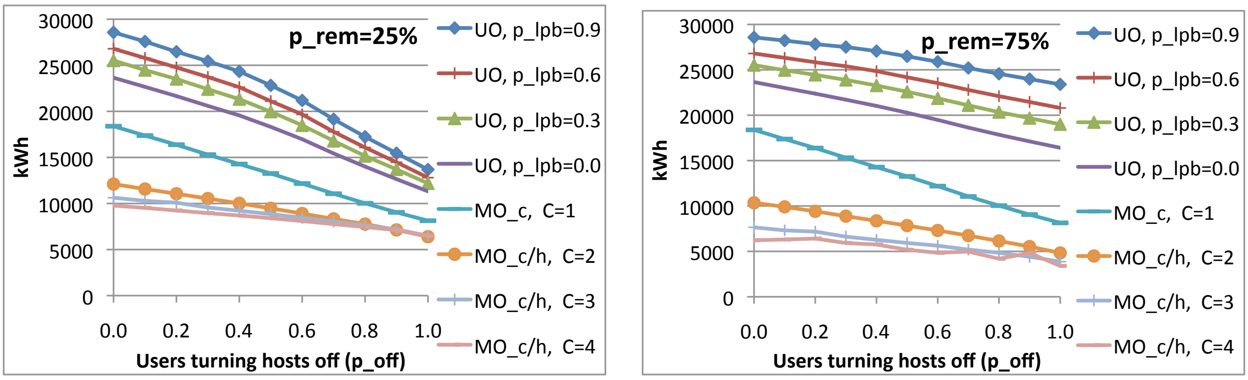

(a) $p_{rem=0.25}$ (b) $p_{rem=0.75}$

Figure 5: Energy consumption of office hosts

the more users show energy-efficient behaviour (x-axis). The lower 4 curves show the energy consumption of an 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 (as explained in 6). The curve illustrates energy savings achieved by automatically suspending unused PDEs (together with their hosts). It can be observed that by just applying this management a significant reduction of energy is achieved. When the consolidation factor is raised to $C = 2$, additionally consolidation is done in the system and PDEs that are not used locally are moved to other hosts. Again, a significant portion of energy can be saved by this method in the MO. It can be observed that a further increase of the consolidation factor ($C = 3$ or $C = 4$) reduces the consumption of energy, but in smaller portions. $C = 3$ or $C = 4$ show more effect, if the ratio of remote users is higher (Figure 5 (b)), because especially the PDEs of remote users are easy to consolidate. The results of MO_c and MO_h are nearly the same for $C = 2, 3, 4$ (not distinguishable in the graph) because the high number of users lessens the differences.

Energy savings of MOs are also illustrated in Figure 6 for $C = 2$ and $C = 3$. In this simulation all parameters are set similar to previous simulations. The average percentage of users that turn their host off (energy-efficient users) is set to 40% and the average percentage of non functional low-power modes is set to 60%. The x-axis shows a growing number of users in the office and the y-axis shows the energy savings that are achieved by the MO approaches. The savings are highly dependent on the user behaviour in the office. It can be observed, that the highest ratio of remote users leads to highest energy savings. It can be clearly observed in Figure 6 that the suggested MO saves up to 67% of energy with $C = 2$ and up to 75% of

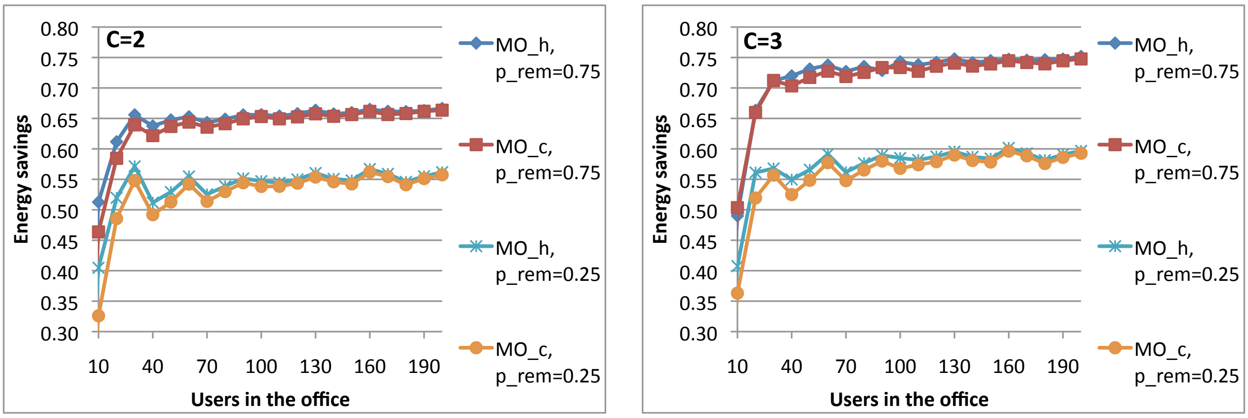

(a) $C = 2$ (b) $C = 3$

Figure 6: Energy savings with an increasing number of users

energy with $C = 3$. Even offices with low remote usage (25%) or offices with a small number of employees (e.g., 20 to 30) significantly save energy with both consolidation factors.

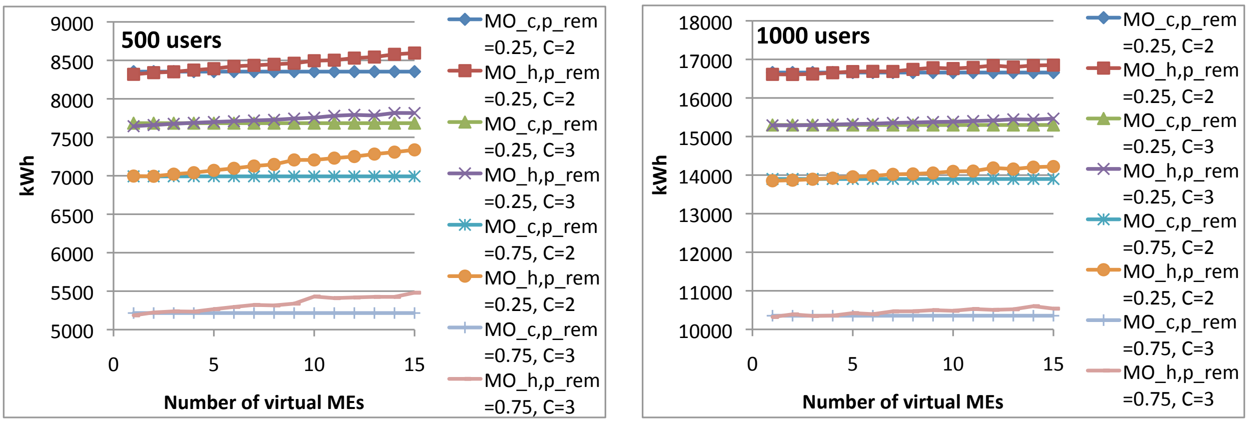

(a) 500 users (b) 1000 users

Figure 7: Virtual MEs (MO_h) versus centralized ME (MO_c)

Advantages of MO_h compared to the MO_c approach are visible in Figure 6 (a). The smaller the number of users is, the more savings can be achieved by the hybrid approach. The highest difference in both approaches appears with 10 users (or less). Another difference between the two approaches – the difference of their energy consumption – is depicted in Figure 7. 500 users (a) and 1000 users (b) have been simulated over a time period of 4 weeks. The x-axis shows the number of virtual MEs in the MO_h approach. The y-axis shows the energy consumption of hosts in the office for both approaches. The energy usage in the

MO_h approach increases with each added ME, whereas the usage of the MO_c approach is constant (a single centralized server). It can be observed that the energy consumption only increases slowly with an increasing number of virtual MEs. The figures illustrate that several virtual MEs in the MO_h approach can be operated with similar energy costs as a single centralized ME. Up to 4 virtual MEs can be operated in the presented examples (with 1000 users and $p_{rem=0.25}$) causing the same energy costs as a dedicated server. A higher number of MEs increases scalability resilience of the managed office environment (see Section 4).

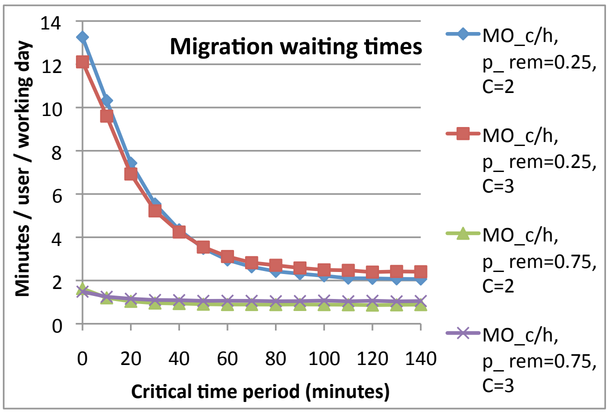

(a) Migration waiting times

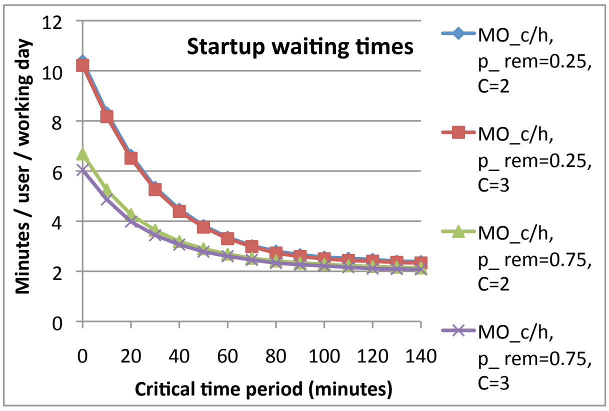

(b) Startup waiting times

Figure 8: Migration and startup times (per user and per working day)

Figure 8 compares the UO and the MO in terms of the service which is provided to users. In the simulation all parameters are set similar to previous simulation. The number of users is set to 200 and the simulation time to 12 weeks. The x-axis denotes the critical time period T_c after which a PDE is suspended or migrated. The y-axis illustrates waiting times of users in the MO. Figure 8 (a) illustrates waiting times caused by migrations – users have to wait until the migration has finished to access their PDE locally. Effects on the user during the migration time depend on the implemented type of migration (see Section 3). An implementation of the described transient thin-client mode would lead to a slightly decreased quality of service during this time (e.g., slower performance). Figure 8 (b) illustrates waiting times according to startup of PDEs and hosts, which is more important to the user's quality of service. During these times, the user has actually to wait for the service to be provided. It can be observed that in both graphs 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 40 to 60 minutes, the waiting times per user and day reduce to reasonable values.

The startup times for $p_{rem} = 0.25$, e.g., reduces to less than 4 minutes per day and user, when T_c is set to 45 minutes. This seems to an acceptable trade-off, considering the achieved energy savings. Waiting times may be further reduced by implementing a more comprehensive management algorithm in the system in future work. Further evaluations of the managed office environment are described in [9] and [5].

6. Related Work

The development of energy efficient IT equipment is fostered by labels such as the US Energy Star[9] or the European TCO Certification[10] which rate IT products (mostly monitors) according to their environmental impact. Novel emerging 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 [17], 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 [18]. 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 [19] 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 that is managed individually by a user. In contrast, the methods described in this paper focus 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[11], Adaptiva Companion[12], FaronicsCore[13], KBOX[14], or LANrev [15]. 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.,

[9] http://www.eu-energystar.org
[10] http://www.tcodevelopment.com
[11] http://entisp.com/pages/eiPowerSaver.php
[12] http://www.adaptiva.com/products_companion.html
[13] http://faronics.com/html/CoreConsole.asp
[14] http://www.kace.com/solutions/power-management.php
[15] http://www.lanrev.com/solutions/power-management.html

backup processes) and for remote usage. Such solutions, however, rely on the capability of the host to switch to low-power modes which 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 the suspension. 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.

Thin-client/terminal-server approaches use data-centre technology to provide energy-efficient services in office environments. User environments (similar to PDEs) are provided by terminal servers and users can access these environments via energy-efficient thin clients. Common terminal-server software products are Citrix XenApp[16], Microsoft Windows Server 2008[17], or the Linux Terminal Server Project[18]. Similar to the approach suggested in this paper, such approaches foster a resource sharing among users in the office environment. However, this approach is based on the usage of additional hardware in the office (energy-efficient thin clients and terminal servers) and PDEs are provided in a centralized way by the terminal server. Instead, the approach suggested in this paper utilizes available hosts in office environments and shares resources among them.

The term cloud computing [20, 21] has been introduced recently and refers to data-centre-based services, stored in ubiquitous computing clouds and is strongly related to grid computing [22]. 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

[16]http://www.citrix.com/XenApp
[17]http://www.microsoft.com/windowsserver2008
[18]http://www.ltsp.org

homogeneous hardware which is located close to each other in racks, interconnected via 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.

In [23, 24, 25] 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 moved, 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.

7. Conclusions and Future Work

This paper has presented an architecture that manages resources in office environments in energy efficient ways. A shift from current decentralized resource management approaches (per user) to a centralized resource management approach (per office) is suggested . The proposed solution extends available power-management approaches and is opposed to data-center based thin-client/terminal-server solutions. It exploits available energy savings in office environments by managing office resources based on the behaviour of users. Resource virtualization technologies (system virtualization and peer-to-peer overlays) are used to suspend idle services and to consolidate underutilized services on a small number of hosts. Furthermore, this paper has discussed and compared different peer-to-peer approaches that can be applied to achieve an energy efficient management. Simulation results indicate that 75% energy savings can be achieved in certain user scenarios, independent of the used peer-to-peer approach.

In future work, the suggested architecture will be refined (e.g., in terms of consolidation algorithms, resilience and security issues), further user scenarios will be analyzed, and different office environments (e.g., based on portable computers) will be compared with each other.

8. Acknowledgements

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 framework programme in the context of the EuroNF Network of Excellence (grant agreement no. 216366) and the P2P-Next project (grant agreement no. 216217).

9. References

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

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

[3] C. Webber, J. Roberson, M. McWhinney, R. Brown, M. Pinckard, J. Busch, After-hours power status of office equipment in the usa, Energy-the International Journal 31 (14) (2006) 2487–2502.

[4] 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.

[5] A. Berl, H. de Meer, A Virtualized Energy-Efficient Office Environment, in: Proceedings of International Conference On Energy-Efficient Computing and Networking (e-Energy 2010), Passau, Germany, April 13-15, 2010, ACM digital library and IEEE Computer Society Press, 2010.

[6] F. Bellard, QEMU, a fast and portable dynamic translator, in: Proceedings of the USENIX Annual Technical Conference, FREENIX Track, 2005, pp. 41–46.

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

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

[9] A. Berl, H. de Meer, An Energy-Efficient Distributed Office Environment, in: Proc. of the 1st Int'l Conf. on Emerging Network Intelligence EMERGING 2009 - Includes the 5th European Conf. on Universal Multiservice Networks (ECUMN 2009), IEEE Computer Society, 2009, pp. 117–122.

[10] 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 (2007) 16–25.

[11] R. Steinmetz, K. Wehrle, Peer-to-Peer Systems and Applications (Lecture Notes in Computer Science), Springer-Verlag New York, Secaucus, NJ, USA, 2005.

[12] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, H. Balakrishnan, Chord: A scalable peer-to-peer lookup service for internet applications, in: SIGCOMM '01: Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications, ACM Press, 2001, pp. 149–160.

[13] I. Gupta, K. Birman, P. Linga, A. Demers, R. V. Renesse, Renesse. kelips: Building an effcient and stable p2p dht through increased memory and background overhead, in: In Proceedings of the 2nd International Workshop on Peer-to-Peer Systems (IPTPS '03), 2003.

[14] J. Falkner, M. Piatek, J. P. John, A. Krishnamurthy, T. Anderson, Profiling a million user dht, in: IMC '07: Proceedings of the 7th ACM SIGCOMM conference on Internet measurement, ACM, New York, NY, USA, 2007, pp. 129–134.

[15] B. Beverly Yang, H. Garcia-Molina, Designing a super-peer network, in: Data Engineering, 2003. Proceedings. 19th International Conference on, 2003, pp. 49–60.

[16] R. Atkinson, S. Bhatti, S. Hailes, Mobility through naming: impact on dns, in: MobiArch '08: Proceedings of the 3rd international workshop on Mobility in the evolving internet architecture, ACM, New York, NY, USA, 2008, pp. 7–12.

[17] Intel, Intel white paper 30057701. wireless intel speedstep power manager: Optimizing power consumption for the intel pxa27x processor family (2004).

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

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

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

[21] A. Berl, E. Gelenbe, M. di Girolamo, G. Giuliani, H. de Meer, M. Q. Dang, K. Pentikousis, Energy-Efficient Cloud Computing, The Computer Journal Advance Access (2009) 1–7.

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

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

[24] A. Berl, H. de Meer, H. Hlavacs, T. Treutner, Virtualization Methods in Future Home Environments, IEEE Communications Magazine 47 (12) (2009) 62–67.

[25] A. Berl, H. Hlavacs, R. Weidlich, M. Schrank, H. de Meer, Network Virtualization in Future Home Environments, in: Proc.

of the 20th Int'l Workshop on Distributed Systems: Operations and Management (DSOM09), Vol. 5841 of Lecture Notes in Computer Science (LNCS), Springer Verlag, 2009, pp. 177–190.