

# Distributed energy efficiency in future home environments

Helmut Hlavacs · Roman Weidlich ·  
Karin A. Hummel · Amine M. Houyou ·  
Andreas Berl · Hermann de Meer

Received: 1 December 2007 / Accepted: 20 June 2008 / Published online: 1 August 2008  
© Institut TELECOM and Springer-Verlag France 2008

**Abstract** In this paper, a new architecture for sharing resources among home environments is proposed. Our approach goes far beyond traditional systems for distributed virtualization, like PlanetLab or grid computing, as it relies on complete decentralization in a peer-to-peer (P2P) like manner and, above all, aims at energy efficiency. Energy metrics are defined, which have to be optimized by the system. The system itself uses virtualization to transparently move tasks from one home to another to optimally utilize the existing computing power. We present an overview of our proposed architecture, consisting of a middleware interconnecting computers and routers in possibly millions of homes using P2P techniques. For demonstrating the potential energy saving of distributed applications, we present an analytical model for sharing downloads, which is verified by discrete event simulation. The model represents an optimistic case without P2P over-

head and fairness. The model allows to assess the upper limit of the saving potential. An enhanced version of the simulation model also shows the effect of fairness. The fairer the system gets, the less efficient it is.

**Keywords** Energy efficiency · Home networks · Peer-to-peer · Modeling · Simulation

## 1 Introduction and motivation

Modern home environments are envisioned as multimedia homes consisting of a multitude of networked devices presenting and managing multimedia services, like video streaming, IP telephony (VoIP), video content delivery, and enabling remote access to home services. Examples for platforms supporting these services are OSGI<sup>1</sup> and UPnP.<sup>2</sup>

Although most of the mentioned services are already available today, future home environments are facing new challenges. On the one hand, a shift from multi-service networks toward multi-network services is likely to be seen. These services are no longer built by one network provider and a single access network, but rather run on multiple networks supported by several providers. On the other hand, more and more *always-on* services are requested by home users. Always-on services include, for instance, file-sharing or other peer-to-peer (P2P) services, multimedia streaming, or remote home monitoring/control. Furthermore, ubiq-

---

H. Hlavacs (✉) · R. Weidlich · K. A. Hummel  
Department of Distributed and Multimedia Systems,  
University of Vienna, Lenaug. 2/8, 1080 Vienna, Austria  
e-mail: helmut.hlavacs@univie.ac.at

R. Weidlich  
e-mail: roman.weidlich@univie.ac.at

K. A. Hummel  
e-mail: karin.hummel@univie.ac.at

A. M. Houyou · A. Berl · H. de Meer  
Faculty of Computer Science and Mathematics,  
University of Passau, Innstr. 43, 94032 Passau, Germany  
e-mail: houyou@fim.uni-passau.de

A. Berl  
e-mail: berl@fim.uni-passau.de

H. de Meer  
e-mail: demeer@fim.uni-passau.de

<sup>1</sup><http://www.osgi.org/>.

<sup>2</sup><http://www.upnp.org/>.

uitous computing technology [15], like smart artifacts consisting of sensors and actuators, are integrated into future homes to support home automation services. In this paper, computing is shifted beyond human awareness involving sensing technologies like measuring environmental phenomena (e.g., temperature, humidity, etc.) or motion detection and position recognition (e.g., to support location-based services in the home). Recent research projects investigate the potential of future home environments, like Amigo<sup>3</sup> (Ambient Intelligence for the Networked Home Environment), the Place Lab,<sup>4</sup> and Easy Living.<sup>5</sup>

These current and future always-on services rely on home computers running on a 24/7 basis while, most probably, being not fully utilized. Always-on computers consume considerable amounts of energy worldwide. For example, a low-cost PC consumes about 100 W if switched on, a multimedia PC consumes 148 W, and only a few watts are consumed if the computer is hibernating.<sup>6</sup> In addition to increased CO<sub>2</sub> balance caused by high energy consumption, energy consumption is seen as major cost factor for servers [5], which is becoming true also for home networks.

This paper discusses the current state of the research project virtual home environments (VHE), interconnecting the Universities of Vienna, Passau, and Cantabria and being sponsored by the Network of Excellence Euro-FGI.<sup>7</sup> VHE proposes a distributed approach to assure energy efficiency for future home networks by means of resource sharing; that is, home services either run locally or they are executed on a remote connected home network. In this paper, resource sharing allows to shift home services (load) to other under-utilized home networks and, thus, allows to put some computers into hibernate mode. The approach of *distributed energy efficiency* is based on home network virtualization, which supports remote execution in virtual machines, a P2P overlay utilized for distributed management, and a distributed algorithm for deciding where to execute home services most efficiently and which home networks should be contributing resources.

One of the novelties of the approach is the distributed energy saving aspect that has not been addressed so far. The implied reduced CO<sub>2</sub> emission is not quan-

tified but is assumed to result from the new system. The second novelty is the interconnection of home networks in a robust, scalable manner to share resources and energy. Our approach is related to other work done in the area of distributed resource sharing, as PlanetLab or Grids, which often lack full decentralization, and to energy efficiency research that so far concentrates on local energy saving, including data centers (see Section 2). While optimizing for energy efficiency, important system characteristics are considered as well in terms of availability, security, fairness (i.e., contributing and retrieving equal amounts of energy), and quality of service (QoS) in particular necessary for the multimedia and home automation services. We propose distributed monitoring of energy and performance metrics (generation of statistics) and distributed decision making. These functions are implemented in a distributed management component utilizing a P2P overlay. Hereby, virtualization of home networks enables the distributed approach by supporting the shifting of home services (see Section 3). A novel system architecture is proposed and described, which details the components necessary for interconnecting home networks (see Section 4). Finally, a discussion of the potential for energy saving in such a distributed environment is provided based on analytical performance evaluation applied to sharing downloads (see Section 5). Section 7 concludes the paper.

## 2 Related work

The most comparable platform to the architecture proposed in this paper is PlanetLab [1, 11]. Although PlanetLab represented a similar vision of an open distributed platform for developing large distributed applications, it stagnated at the stage where it has become an experimental platform for testing large Internet-based research. This paper proposes an extension to the vision of PlanetLab focused on a distributed home environment and puts energy sharing as a focal point, which would incite users to offer their home PCs for resource sharing.

Similar to server environments [5], energy consumption is becoming a major problem in home networking, as energy costs tend to exceed that of hardware. Koomey [10] mentions that today's energy consumption of volume, mid-range, and high-end servers in the USA and worldwide has doubled over the period from 2000 to 2005. The total power demand in 2005 (including associated infrastructure) is equivalent to about five 1,000 MW power plants for the USA and 14 such power plants for the world [10].

<sup>3</sup><http://www.hitech-projects.com/euprojects/amigo/>.

<sup>4</sup>[http://architecture.mit.edu/house\\_n/placelab.html](http://architecture.mit.edu/house_n/placelab.html).

<sup>5</sup><http://research.microsoft.com/easyliving/>.

<sup>6</sup>Energy Star Europe calculator: [http://www.eu-energystar.org/en/en\\_007c.shtml](http://www.eu-energystar.org/en/en_007c.shtml).

<sup>7</sup>[http://eurongi.enst.fr/en\\_accueil.html](http://eurongi.enst.fr/en_accueil.html).

Nevertheless, energy efficient computing is not a new topic. With the need of a longer battery life time in laptops, for instance, several techniques such as Speed-Step [8], PowerNow, Cool'nQuiet, or demand-based switching [17] have been developed as local power-saving measures. These measures enable slowing down the clock speeds (Clock Gating) or powering off parts of the chips (Power Gating), if they are idle [6, 16]. A further power adaptive technique is based on sensing whether the computer has been left idle, based on human-machine interaction input components (e.g., keyboard, mouse, touch-pad, etc). The longer the computer is left idle, the more hardware elements are turned off or suspended, while allowing a turn on mechanism without loss of state or information. This mechanism allows a gradual reduction of power usage. However, turning hardware off does not always imply that a computing system is energy efficient. Energy efficiency can be measured in performance per watt [9]. One way to attain a better performance per watt has been achieved by using virtualization technology. Virtualization can be seen as splitting an underlying hardware entity into smaller identical virtual entities running isolated from each other. In data centers for instance, the rack-mounted servers were configured to run a single workload to guarantee reliability, availability, and scalability of the service. This came at the cost of under-utilized energy expensive machines, which had an average load of about 10% [5]. With virtualization, a virtual machine is dedicated to each service, but can run transparently on any available system next to several other virtual machines. This effective consolidation of servers, i.e., running a machine at a higher utilization [9], is usually done by a central management mechanism. A further energy saving method, which is currently investigated within the context of data centers [13], consists of turning parts of the machines off while taking cooling cost into account [7, 10].

A similar management type of a virtual environment could be found in grids, e.g., Condor [14]. Condor is a workload management system that allows users to submit their jobs to a single queue. The management system distributes the jobs transparently among the computing grid. This functionality, however, is centralized and does not take energy efficiency into account.

It is such a management mechanism and dynamic behavior, which is missing in a platform like PlanetLab. There, virtual environments for users are created centrally, one virtual environment on each PlanetLab machine. However, in PlanetLab [1], shifting load is not trivial; consolidating machines to run at a higher load is not yet possible. Furthermore, there is no automatization in allocating virtual resources to a given

user or to a special application. In our architecture, we aim at the automatic allocation of virtual resources in a distributed environment while consolidating the future home PC and switching those PCs off that run at a light load.

### 3 Distributed energy efficiency

The concept for distributed energy efficiency relies on the concept of interfering characteristics that decide upon where to run home services (where to shift the load to). Hereby, the management algorithm assures fairness, availability, QoS, and security while optimizing for energy saving and energy efficiency. *Fairness* means that each home should consume approximately as much as it contributes to the system. For reasons of robustness and scaling, a distributed solution is proposed, considering each of these characteristics. The distributed solution is supported by virtualization techniques.

#### 3.1 Energy efficiency optimization and constraints

The goal of the presented approach is to globally minimize the energy consumption, thus maximizing the global energy efficiency. For a number of  $N$  different homes  $h_i$ ,  $1 \leq i \leq N$  the basic energy consumption  $E(T)$  over system time  $T$  is given as:

$$E(T) = \sum_{i=1}^N \int_0^T P_{h_i}(t) dt \quad [\text{joule (or kWh)}],$$

where  $P_{h_i}(t)$  is the power consumed by a home  $h_i$  in watt. In absence of measurement possibilities of homes, the energy consumption of a home might as well be estimated by assigning an energy class level to the home.

To calculate the energy efficiency, the workload introduced by the home network services is related to energy consumption; thus, the work carried out by all homes is defined as:

$$L(T) = \sum_{i=1}^N \int_0^T L_{h_i}(t) dt,$$

where  $L_{h_i}(t)$  describes the work caused by the home services at time  $t$  (seen as the work *output* of a home). Similar to [12], we define the overall energy efficiency of the system to be maximized by

$$\eta(T) = \frac{L(T)}{E(T)}, \quad (1)$$

where it is assumed that  $E(T) \neq 0$  kWh. If the energy consumption can be reduced by sharing, the energy efficiency will increase.

Additionally, the system assures a certain degree of trust in the non-functional characteristics of home services, thus requiring more computing power, which, as a consequence, causes additional energy consumption. The addressed characteristics are *availability*, *security*, *fairness*, and *QoS*, which are constraints to the optimization problem to minimize energy consumption and to maximize energy efficiency.

Based on these basic energy formulas, a distributed solution is proposed, where load, i.e., home services, are shifted between homes to optimize  $E(T)$  and  $\eta(T)$  (to be more precise, a combination of both optimization problems). In the absence of a central management, the global behavior emerges based on the local behavior of homes. Each home conducts performance measurements and monitoring of energy consumption as well as a decision algorithm to determine whether to provide resources for home services. In addition to energy consumption, for example, the mean time to failure and the mean time to repair are calculated to describe availability, the mean load caused by home services are monitored for reasons of fairness, and the mean data transfer rates for up and down links address QoS constraints. For security reasons, mutual monitoring of past malicious behavior is performed resulting in security levels assigned to homes.

The distributed optimization algorithm for decision taking is based on building groups of homes that exchange performance, security, and energy status information to build partial views of the global state. Based on this information, ideally, each home can execute an identical algorithm, deciding upon the home's contribution to solving the optimization problems by converging toward the optimal energy saving or energy efficiency while considering the services' requirements.

### 3.2 Decentralized virtualization

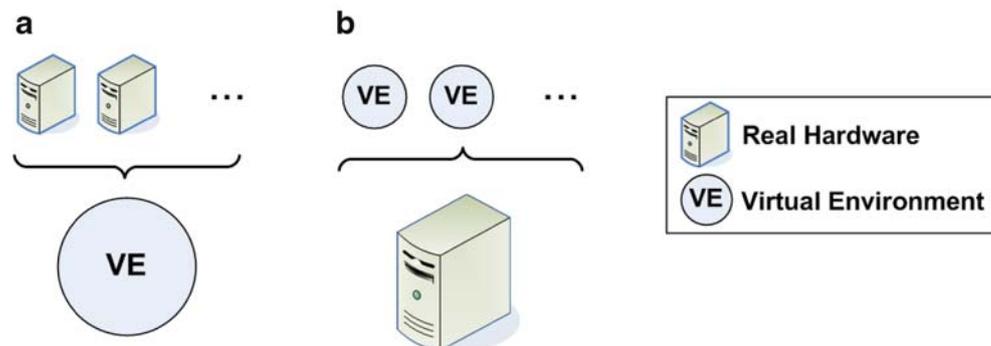
Mechanisms for resource virtualization have been used in different contexts, aiming at different results. Three examples (grid computing, server virtualization, and virtualization in PlanetLab) are described to clarify their different targets and to illustrate the next step taken by the architecture that is proposed in this paper.

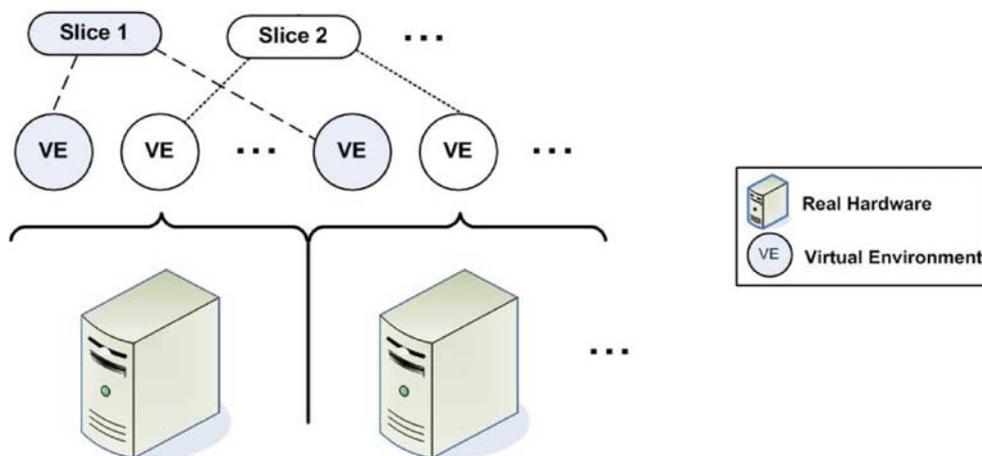
In grid and cluster computing (e.g., in Linux clusters), virtualization is used to aggregate a pool of hardware resources. In this context, virtualization aims at hiding the complexity of aggregating several machines in a grid/cluster from the user.

The user accesses the aggregated hardware (e.g., high number of CPUs, large amount of memory) as a single virtual environment (e.g., a single Linux shell). This kind of virtualization is shown in Fig. 1a. A number of real machines are aggregated to a single virtual environment (aka *virtual organization*). In contrast to the compositional grid virtualization, server virtualization uses virtualization methods in a segmenting manner. Server virtualization aims at splitting hardware resources into several smaller virtual environments, enabling more than one virtual environment on a single hardware. Servers are virtualized to achieve load balancing, to increase resilience, and to save hardware/energy by consolidation, e.g., in data centers. In Fig. 1b, this kind of resource virtualization is shown. A single hardware is split into several virtual environments (aka *virtual machines*).

PlanetLab [1] faces a more complex, distributed scenario of virtualization. Hardware resources are spread all over the planet, interconnected via the Internet, without the use of special high-performance links. Within PlanetLab, every single machine is split into virtual environments similar to server virtualization. These virtual environments are organized in slices. More precisely, a slice is represented by one single

**Fig. 1** a Grid virtualization and b server virtualization



**Fig. 2** Virtualization in PlanetLab

virtual environment per available PlanetLab machine. Thus, a user who has booked a slice receives one Linux shell per PlanetLab machine. This scenario is illustrated in Fig. 2. It is important to see that virtual environments in a slice are not aggregated like resources in a grid. No further abstraction than one shell per machine is provided, leaving users with the problem of dealing with dozens or even hundreds of shells simultaneously.

This paper proposes a distributed virtualization solution that goes one step further than the virtualization in PlanetLab. An architecture is suggested, in which slices are variable in size (number of involved virtual environments) and change their location dynamically. These extended slices are called *flexible slices*. As an example, a flexible slice might consist of four virtual environments, which are located in the current home network at one time, and consist of seven virtual environments, which are located in other home networks at another time. However, similar to the virtualization used in grid/cluster computing, this complexity is hidden from the user. The user experiences a single virtual environment (virtual organization) in which the resources of the flexible slice are aggregated.

### 3.3 Decentralized management

To take advantage of virtualization, management of the virtualized hardware has to be done. In grids, available resources have to be adequately allocated. In data centers, virtual servers have to be moved, copied, created, and deleted, e.g., for load balancing or consolidation. Similar to the resources of grids, server hardware is usually located close to each other, e.g., in racks or data centers, and interconnected with high-bandwidth links. Therefore, the management of virtualization in grids

and data centers is mainly implemented in a centralized way, where a central management element allocates resources. The *VMWare Infrastructure 3*,<sup>8</sup> for instance, provides such a centralized management element to manage virtual machines in data centers. Although virtualization itself is highly distributed in PlanetLab, the management of hardware and slices is rather centralized. Slices are created, allocated, and managed via a central server. Furthermore, the user of a slice is a central point of management, having to cope with hundreds of virtual machines.

In the approach proposed in this paper, flexible slices have to be managed to provide the envisioned future home environment. Home networks are interconnected by a P2P overlay and share their resources to enable distributed energy efficiency. Always-on services are wrapped into flexible slices (transparent for the users), making them movable within interconnected home environments. Energy saving is achieved by increasing the load on some computers while turning off others. The constraints (fairness, security, availability, and QoS) described in Section 4 in more detail have to be considered within the management decisions. The decision process is based on distributed statistics, which are gathered in the home networks. To achieve a scalable management in a dynamic and vast environment and to avoid single points of failures, the management is decentralized as far as possible. Homes with active computers are involved in the decision process, which concerns all of the interconnected home networks.

<sup>8</sup>[http://www.vmware.com/pdf/vi\\_brochure.pdf](http://www.vmware.com/pdf/vi_brochure.pdf).

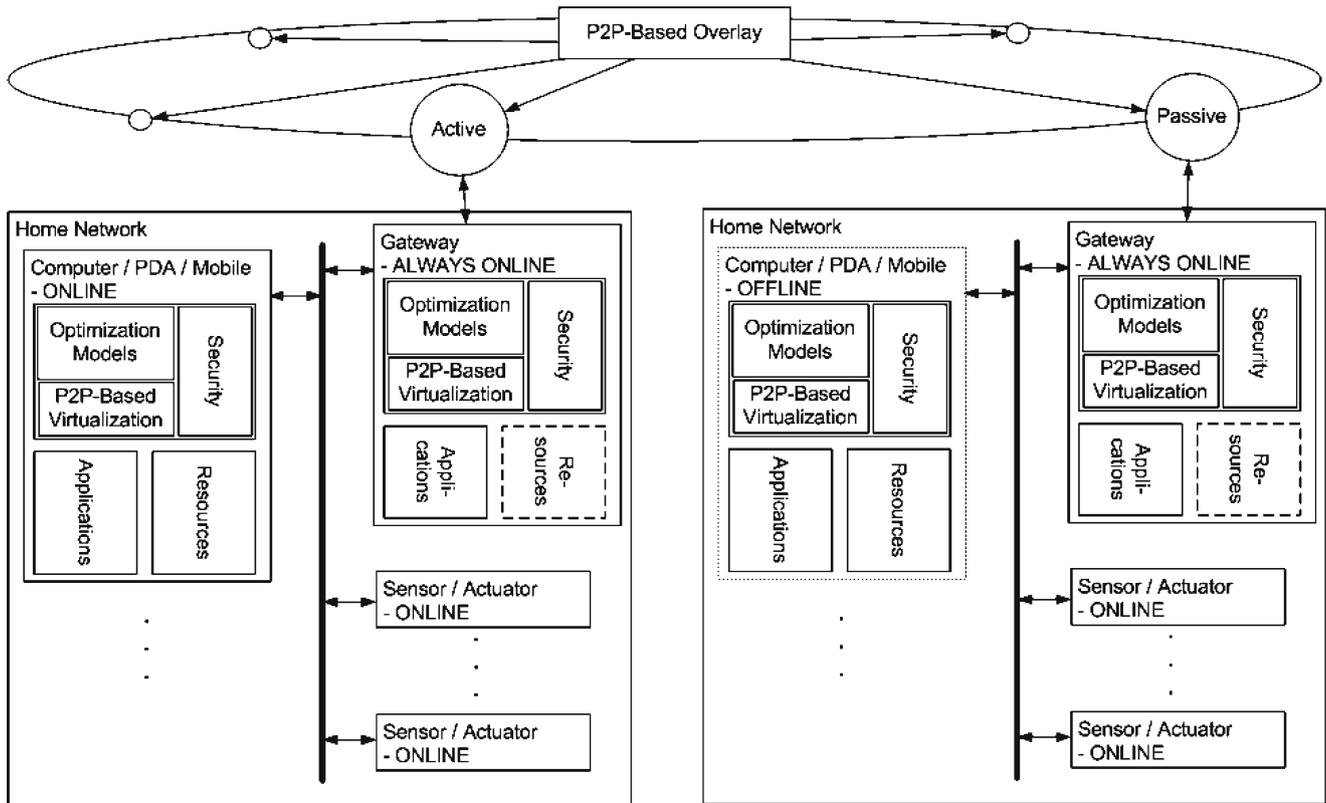


Fig. 3 Distributed energy efficient home network architecture

4 System architecture

The proposed architecture for the distributed energy efficient resource sharing approach consists of interconnected homes. Each *home* is an abstraction from a home network consisting of an always-on gateway (or router) that connects the home network to the Internet, one or several computers and displays, connected peripherals, and sensors and actuators. For interconnection, the homes are using a distributed hash table (DHT) based P2P overlay. Figure 3 shows the proposed architecture. The home network (depicted as a bus system) consists of multi-networks, for example, wireless networks (like WLAN IEEE 802.11g) and wired networks, like serial line connections or Ethernet (for connecting sensors), and a high-speed up-/downlink to the Internet.

Each component of the home, which we refer to as a *node* (e.g., any computer, sensor, actuator, PDA, etc.) is represented by static (like the processor speed and main memory size) and dynamic (like the utilization and the energy consumption) characteristics. Additionally, each node is in one of the states: *active* (online and contributing), *active-blocked* (online but not contributing), or *passive* (in suspended, hibernating, low power

mode). The active-blocked state has been introduced to support the user who wants to stay in control of his/her home equipment. For example, if the user wants to join a massively multiplayer online role playing game, bandwidth and computing power should not be contributed for energy efficient resource sharing, because otherwise the gaming experience might be negatively influenced. Similarly, a home is said to be active if it contributes to the system, active-blocked if it is not ready to contribute, and passive if all possible contributing resources of the home are hibernating.<sup>9</sup> The abstraction of the nodes in terms of their characteristics is aggregated to describe the characteristics of a home.

Virtualization techniques are applied in two ways as described in Section 3. First, the system appears as an abstract *virtual organization* (in compliance with grid technology) to the service that is executed transparently on participating nodes of the system (residing in homes). Second, load distribution and shifting

<sup>9</sup>Note that, in case a home is passive or active-blocked, the gateway is still up, and the home might consume services like home automation services from the distributed home environment.

are implemented by utilizing the technology of *virtual machines*. The management of the load distribution is done in a distributed manner by executing a distributed algorithm on each node. In Fig. 3, the depicted security, optimization model, and P2P-based virtualization modules implement the management functionality and are described below in more detail.

Through virtualization, applications can transparently allocate resources like disk space, CPU time, or bandwidth without knowledge about location or configuration of remote computers by logically separating application modules into *frontends* and *backends*. The frontend implements only few functions of the application like the user interfaces, while the backend implements the heavy-loaded business logic. Due to virtualization, many backends could be assigned to the same frontend, while the distributed execution is hidden. The user only must start his frontend instead of starting the applications on his own computer.

Figure 3 shows that the intelligence of the distributed management layer is situated in each contributing node, which may be both a full-blown PC with large computational resources (but also large energy consumption) or the home router/gateway, which is assumed to be a simple Linux-based diskless computer with small energy needs. Although this gateway is not able to contribute its own resources to be used by other homes, its computational power should be sufficient to maintain a permanent entry in the system wide DHT for representing its particular home. As the gateways are assumed to run permanently (as usually all routers/gateways do), the churn as experienced by the P2P system is thus almost zero.

The management layer is based on three major building blocks. The P2P-based virtualization manages the overlay and provides services like identifying other peers, providing a system-wide distributed database for storing node statistics persistently (including descriptions of the node resource capabilities, energy class, up/downlink capacity, resources contributed so far to the system, etc.), or transferring resource requests from one computer to another.

Above it, optimization models implement the true intelligence of the system. They can be roughly divided into the following submodels:

- Energy efficiency. Once a frontend requests to use the resources of a remote computer, depending on the type of request, this submodel tries to identify a set of nodes that should be selected because selecting them would minimize the global energy consumption and maximize global energy efficiency.

- Fairness. This model uses statistics about how much each home has contributed to the system recently. Given a resource request and a set of nodes (from the energy efficiency model), this model identifies those nodes who should be assigned because they have not contributed much recently.
- Availability. This submodel decides how the service should be replicated. For instance, storing data for other computers or remote home management should be done by using replication to increase availability.
- Privacy. This model tries to maximize the degree of privacy that a service is experiencing. Consider for instance the case that a remote home manages resources of other homes. To prevent the host computer to find out the identity of the managed home, other homes might function as a proxy chain in between.
- QoS. Depending on the application, given a resource request, this model decides whether a particular node is able to host the requested application. For instance, if the user wants to remotely encode video files, the host computer carrying out the work should actually command a large down- and uplink bandwidth and enough free CPU power. These resources, however, would be used only once. A slower computer on the other hand might be sufficient to receive messages from home management services and answer to them. This particular service then would run for a very long time, thus achieving fairness. A third example for QoS decisions is given by the tradeoff between QoS and privacy. Consider again remote home management. When using *long* proxy chains, the degree of privacy is extremely high, whereas the important QoS parameters latency and bandwidth will be much worse. Thus, for many applications, there is a tradeoff between QoS and privacy.
- Security. This model is part of the P2P layer as well as being part of the optimization models. At the P2P layer, it provides services for encryption and key exchange. At the optimization layer, it mainly governs the distributed voting process. Voting is necessary because malicious nodes may try to create damage in other homes. Consider once more home management. Shutting down heating might be dangerous and cause damage in winter. Thus, such possibly dangerous applications might rely on a majority voting, where for instance the home gateway acts as a policer, and only commands may pass which have been signed by several other homes, rather than by only one.

### 5 Analytical evaluation

To investigate the potential energy saving by cooperation we have developed various analytical models. In this work we focus on a simple download scenario. Other models, e.g., for CPU sharing or home management can be found in [4]. In the download scenario computers may share downloads with each other. As we are only interested into the potential energy saving, security and privacy concerns are not included into the model. Downloads are carried out via a conventional file-sharing tool like *KaZaa*, *eMule* or *BitTorrent* from the Internet, i.e., from computers which are not part of the modelled scenario. A computer *A* may send a download request to another computer *B*, which will then carry out the download. This way, downloads can be shared and only a small number of computers must be active and thus consume energy. Other computers may sleep, thus not consuming energy at all. Once the download on computer *B* has finished, *B* sends back the file to computer *A*, here waking up *A*, which will then again consume energy as long as the transfer is going on. As a simplification we assume that computers being active because they download for others, always download their own files.

Furthermore, it is assumed that downloads do not use the whole downlink bandwidth  $B_d$  as given by the Internet connection. Instead, as is experienced with real life file-sharing tools, the download bandwidth for one single file is limited by some upper limit but, on average, uses  $B_l$  Kbit/s with  $B_l < B_d$ .  $B_l$  usually depends on the number of seeders and on properties of the used file-sharing tool. The scenario is described by the following parameters. Parameter  $N$  denotes the number of computers in the scenario, while  $M = \lfloor B_d/B_l \rfloor$  denotes the number of downloads that may be carried out in parallel by each single computer. For instance, if we assume that a computer's raw downlink bandwidth is  $B_d = 4$  Mbit/s and each download on average consumes  $B_l = 200$  Kbit/s, then  $M = 20$  downloads can be carried out concurrently. Parameter  $\lambda$  denotes the arrival rate of download requests at each single computer,  $F$  denotes the average file size,  $t_l = F/B_l$  denotes the average time it takes for downloading a file, and thus  $\mu = 1/t_l$  denotes the rate at which each download is finished. For instance, if the size of a file on average is  $F = 100$  MBytes and  $B_l = 200$  Kbit/s, then  $\mu = 1/4,000$  downloads finished per second.

To make the model analytically tractable, it is assumed that download requests arrive according to a Poisson process, and download times (and thus file

sizes) are distributed exponentially. The latter assumption is in conflict to the well-known fact that file sizes usually follow a Pareto or lognormal distribution. This will later be accounted for in our future simulations.

We investigate three cases, the local case where no sharing occurs (*local*), the ideal resource sharing case (*ideal*), and the corrected case (*corr*). The two latter cases differ in the way they deal with the actual transfer to the requesting peer: While in the ideal case, this transfer is neglected, in the corrected case, this transfer is included (resulting in additional wake-up time for the requesting computer).

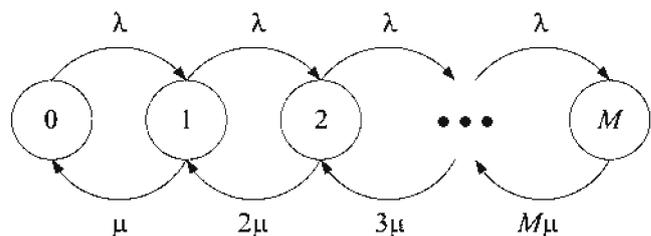
At first, we assume that downloads are carried out on the computer that created the request, i.e., no sharing is going on. Thus, we start by modeling one single computer. The number of downloads carried out by this computer can be modeled by a birth–death process; that is, the process is in state  $k$  if the computer is currently carrying out  $k$  downloads. As  $M$  is the upper bound of downloads, the process has exactly  $M + 1$  states. It is further assumed that, if the process is in state  $M$ , newly generated downloads are lost. This is done as, for the low load investigated here, there is de facto no loss. Otherwise, an  $M/M/M$  queue would be necessary, enforcing also to keep track of the queue length. The process states and transition rates are shown in Fig. 4.

Following [2], simple analysis shows that the probability  $\pi_k$  for being in state  $k$  is given by

$$\pi_k = \pi_0 \frac{1}{k!} \left( \frac{\lambda}{\mu} \right)^k, \quad 1 \leq k \leq M, \quad \text{with}$$

$$\pi_0 = \frac{1}{1 + \sum_{k=1}^M \frac{1}{k!} \left( \frac{\lambda}{\mu} \right)^k}.$$

As  $\pi_0$  denotes the probability that no download is going on,  $1 - \pi_0$  denotes the probability that at least one download is going on, i.e., the computer is active. If there are  $N$  computers, then the expected number



**Fig. 4** Birth–death process for local downloads on one single computer

of active computers  $N_{\text{local}}$  for local downloads only is given by

$$N_{\text{local}} = N \left( 1 - \frac{1}{1 + \sum_{k=1}^M \frac{1}{k!} \left(\frac{\lambda}{\mu}\right)^k} \right). \tag{2}$$

In the next scenario, we assume that computers share downloads; that is, if a computer creates a download request with rate  $\lambda$ , it first searches for an active computer to pass the request to. If there is none, it will start the download itself. Again, the scenario is modeled by a birth–death process, this time by modeling the state of all computers. As there are  $N$  computers and each is able to carry out  $M$  downloads in parallel, in total,  $M \times N$  downloads can simultaneously be carried out; that is, the process has  $M \times N + 1$  states as shown in Fig. 5.

The solution of this process is similar to the one above, yielding

$$\pi_k = \pi_0 \frac{1}{k!} \left(\frac{N\lambda}{\mu}\right)^k, \quad 1 \leq k \leq NM, \quad \text{with}$$

$$\pi_0 = \frac{1}{1 + \sum_{k=1}^{NM} \frac{1}{k!} \left(\frac{N\lambda}{\mu}\right)^k}.$$

When assuming zero communication overhead and not taking into account sending back the download results (*ideal* situation), then the number of active computers necessary to carry out  $k$  downloads is  $a = \lceil k/M \rceil$ . In other words, no computer must be active in state zero,  $a = 1$  computer must be active in the states 1 to  $M$ ,  $a = 2$  for the states  $M + 1$  to  $2M$ , and so on. The probability for needing exactly one active computer is thus given by the sum of the  $\pi_k$ ,  $1 \leq k \leq M$ , and in general, the probability for needing exactly  $a$  active computers is

therefore the sum of the  $\pi_k$ ,  $(a - 1)M + 1 \leq k \leq aM$ . For computing the expectation  $N_{\text{ideal}}$  of  $a$ , we derive

$$N_{\text{ideal}} = \sum_{a=1}^N a \text{P}(a \text{ computers necessary})$$

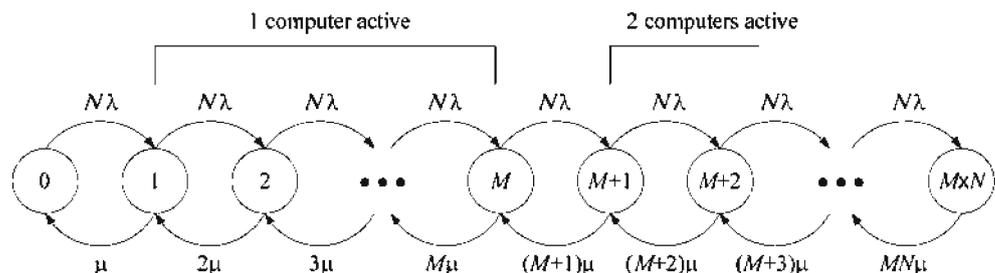
$$= \sum_{a=1}^N a \sum_{k=(a-1)M+1}^{aM} \pi_k. \tag{3}$$

To catch the effect of additional transfer to computer  $A$ , after the download has finished on computer  $B$ , the system is observed for a long time  $T$ . Then, the total time that computers are active within  $T$  is given by  $N_{\text{ideal}}T$ , and the time that the system was in state  $k$  is given by  $\pi_k T$ . From this, it follows that the number of finished downloads while being in state  $k$  is given by  $\pi_k T k \mu$ . As all  $N$  computers contribute equally to the system load, i.e., all create download requests with the same  $\lambda$ , the origins of download requests are distributed evenly among all computers, but only  $\lceil k/M \rceil$  of them are active, while  $N - \lceil k/M \rceil$  are passive. The fraction of finished downloads carried out for currently sleeping computers is thus  $(N - \lceil k/M \rceil)/N$ . It follows that, on average, the number of downloads finished in state  $k$ , which were carried out for a *currently sleeping* computer, is given by

$$\pi_k T k \mu \frac{N - \lceil k/M \rceil}{N}.$$

The time for sending back the result to the initiating computer is given by  $t_u = F/B_u$ , here taking the full raw uplink bandwidth  $B_u$  given by the Internet connection (e.g.,  $B_u = 1$  Mbit/s), which is considered to be much faster than the average download bandwidth  $B_l$  limited by the file-sharing tool. Thus, when sending back a finished download to a computer that was sleeping previously, the sleeping computer must be woken up and must be active for at least  $t_u$  seconds. It follows that, when observing the system for  $T$  seconds, the additional active time  $T_{\text{trans}}$  for transferring back

**Fig. 5** Birth–death process for simultaneous downloads of  $N$  computers



finished downloads to computers that have been sleeping previously is given by

$$T_{\text{trans}} = t_u \sum_{k=1}^{MN} T \pi_k k \mu \frac{N - \lceil k/M \rceil}{N}.$$

When considering additionally that  $t_u = F/B_u$  and  $\mu = B_l/F$ , the number of computers *additionally* running for receiving transfers is then given by

$$N_{\text{trans}} = \frac{T_{\text{trans}}}{T} = \frac{B_l}{B_u} \sum_{k=1}^{MN} k \pi_k \frac{N - \lceil k/M \rceil}{N}. \tag{4}$$

Equation 4 is in accordance with the simple intuition that active time in the sharing scenario is determined by the relation between the download bandwidth  $B_l$  and the raw uplink bandwidth  $B_u$ . The *corrected* average number  $N_{\text{corr}}$  of active computers observed, including the time for downloading and for transferring back, is then defined as

$$N_{\text{corr}} = N_{\text{ideal}} + N_{\text{trans}}. \tag{5}$$

Figure 6 shows results for  $N = 100$  homes,  $F = 100$  MByte,  $B_d = B_u = 8$  Mbit/s, and  $B_l = 1$  Mbit/s, while Fig. 7 shows the same for  $N=1,000$  homes. With arrival rate  $\lambda$ , each single computer generates a certain number of download requests per week, shown at the  $x$ -axis. The possible saving of computer energy is reflected by the difference between the number of active computers in the local case (2) and the corrected case (5). It can be seen that, even when taking into account the distribution overhead, i.e., sending back the files to the requesting computers, the shared scenario (Corr) can save a substantial amount of energy. For instance, in the scenario  $N = 1,000$ , when assuming that each computer consumes 100 W and creates 35 download

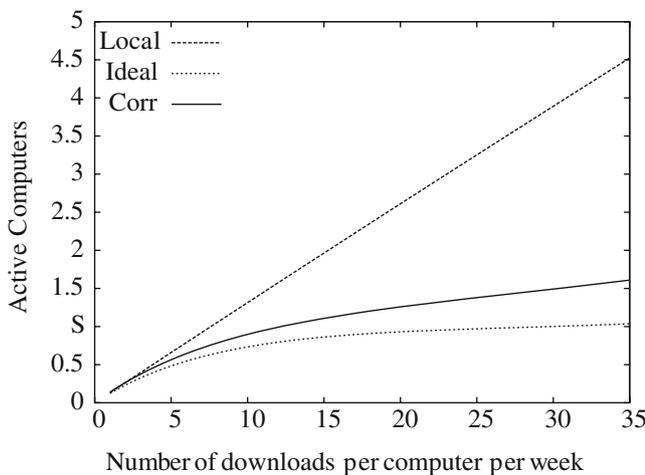


Fig. 6 Number of active computers for  $N = 100$  homes

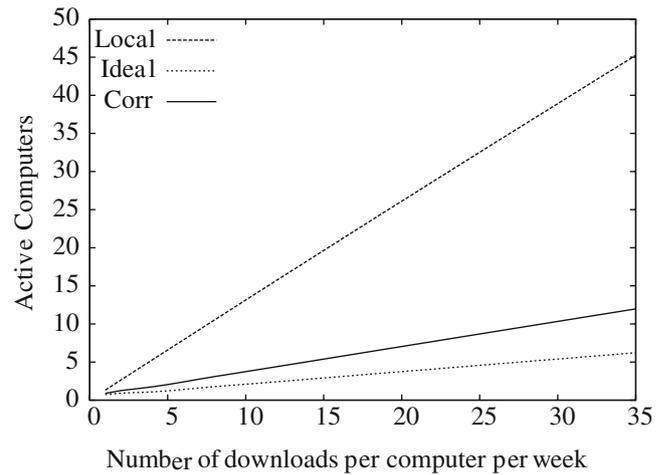


Fig. 7 Number of active computers for  $N = 1,000$  homes

requests every week, without cooperation, 1,000 computers would *constantly* consume more than 4.5 kW on average just for downloading files, while cooperating computers would only consume about 1.2 kW for the same task.

It must be noted that Eq. 5 does not take into account some details, such as fairness, representative file size distribution, and protocol overhead. Additionally, it assumes that, after the last download has finished, the downloading computer immediately may go to sleep, while in reality it must still send the file to the file owner.

### 6 Simulation results

To validate the analytical model and to include more details, we developed a discrete event simulation tool and implemented the download scenario as a case study. The model consists of  $N$  homes interconnected by a central Internet, as shown in Fig. 3. The P2P overlay uses superpeers for keeping track of busy or idle homes.

Computer nodes cycle through the states depicted in Fig. 8. Initially, all computers begin in the passive state, which represents hibernation. In this state, the computer does not consume any energy. According to the download arrival rate  $\lambda$ , each computer wakes up and decides to start one download task. This is represented by a switch to the state active-blocked. At this point, the computer asks its respective superpeer to send a list with currently downloading computers that could take over the task. If no such computer is available, the computer switches to the active state and starts to download the file itself. In this state, the computer also

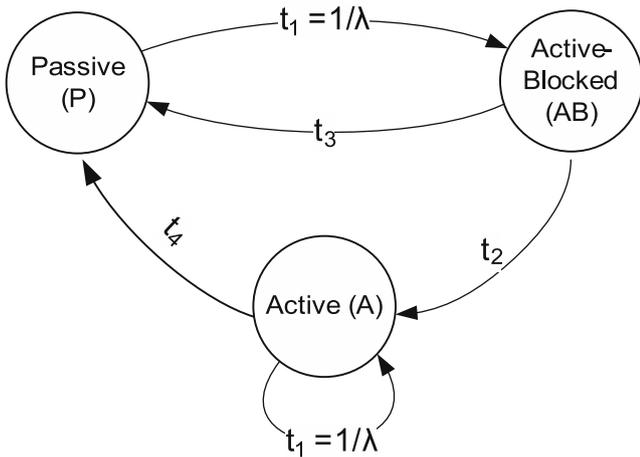


Fig. 8 State transitions of nodes

accepts download requests from other computers. In case an active computer *A* accepts a download task from computer *B*, it starts the respective download and computer *B* may immediately go to the passive state. Upon completion, computer *A* transfers the download back to the owner *B* with full upstream access bandwidth. If, at this point, computer *B* is passive, it is woken up to receive the transfer. Additionally, a node being in state active will continue to start downloads by itself with arrival rate  $\lambda$ .

In an unfair scenario, an active computer might continuously accept tasks and run permanently, thus minimizing the globally consumed energy, but maximizing its locally spent energy. The analytical model given in Eq. 5 represents this case. However, to ensure fairness, an active computer eventually must be allowed to hibernate to stop spending energy and thus money. Therefore, the simulation model also includes a maximum download time  $t_4$ . If an active computer *A* that entered the active state at time  $t_a$ , at time  $t > t_a$  receives a download request from another computer *B*, it estimates the time  $dt$  it will take to finish the download. The request is then accepted by computer *A*, if

$$t + dt \leq t_a + t_4,$$

otherwise it is rejected. Once all downloads are finished and transferred back to their respective owners, an active computer eventually will again enter the passive state.

Figure 9 shows simulation results for  $N = 100$  homes,  $F = 100$  MB,  $B_d = B_u = 50$  Mbit/s,  $B_l = 1$  Mbit/s. Additionally, the pure local case given in Eq. 2 as well as the corrected analytical model Eq. 5 are depicted. The figure shows the influence of the fairness time  $t_4$ .

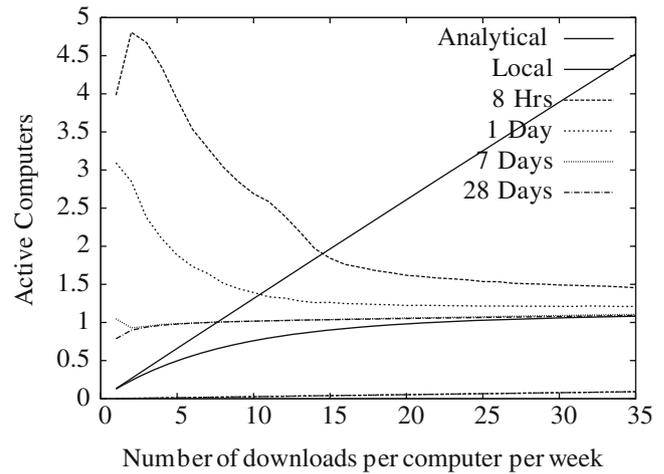


Fig. 9 Mean number of active computers. The solid thick line represents the analytical model (5). The solid thin line represents the pure local case (2). The other thick lines depict simulation results for various values of the maximum download time  $t_4$ . The thin lines seen at the lower right corner show Eq. 4 and the corresponding simulation result

For small values of  $t_4$ , i.e., fairness is achieved quicker, and for small arrival rates, the sharing scenario consumes much more energy than the local scenario. For larger values of  $t_4$  and small arrival rates, the sharing scenario needs much less energy, but still may be worse than the local scenario. Interestingly, there is almost no improvement between  $t_4 = 7$  days and  $t_4 = 28$  days.

The reason for this behavior is due to the rejection of download requests due to the limited maximum download time  $t_4$ . Figure 10 shows the probability that a download request sent to another computer is accepted. For  $t_4 = 8$  h and small  $\lambda$ , this probability is

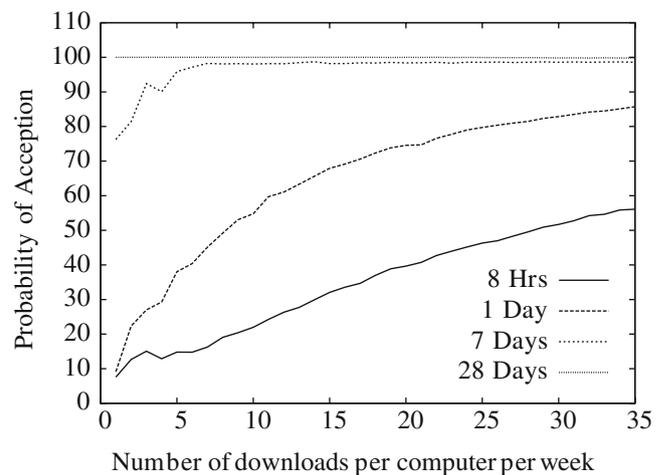


Fig. 10 Probability (percentage) of download acceptance for different values of the maximum download time  $t_4$

around 10%. For larger values of  $t_4$ , this probability rises significantly, and for  $t_4 = 28$  days, it is always close to 100%. However, in this case, the remaining difference between the simulation results and the analytical model is due to the P2P overhead necessary for managing the downloads.

For all values of  $t_4$  and increasing  $\lambda$ , the results approximate the analytical model (5) quite fast. The analytical model therefore gives a good lower bound for the download scenario representing an optimal case without fairness and P2P overhead.

In the lower right corner of Fig. 9, additionally, the number of running owners  $N_{\text{trans}}$  necessary for transferring back the results can be seen. Although not large in value, there are actually two lines, one given by Eq. 4, the other one derived from the simulation runs (in this case for  $t_4 = 1$  day). Both lines cannot be distinguished visually, even when zooming closer. Equation 4 thus indeed provides meaningful values. Note also that the reason for the small value in this scenario can be directly seen in Eq. 4, which depends on  $B_l/B_u$ . While in Figs. 6 and 7,  $B_l/B_u = 1/8$ , in the scenario shown in Fig. 9,  $B_l/B_u = 1/50$ , and the influence of  $N_{\text{trans}}$  is reduced accordingly.

## 7 Conclusion

In this paper, a novel architecture for virtualizing and sharing hardware resources in future home environments is presented. The architecture aims at utilizing existing home resources in such a way that the consumed energy is minimized and the energy is efficiently used. A fully decentralized management system is proposed, interconnecting possibly thousands of homes in a P2P like manner. Energy optimization is done in a decentralized way by converging to a global energy optimum based on energy and performance metrics that have been defined. For the example scenario *file download*, an analytical model is presented that demonstrates the possible amount of energy that can be saved if computers cooperate and share file downloads, rather than if each computer downloads its own files. The model represents an idealized situation without fairness requirements. To assess the influence of fairness, we evaluated the scenario also by using discrete event simulation. For a small system load, the considered fairness method results in a very bad efficiency, even worse than the local case. For increasing load, however, a fair scenario quickly approximates the ideal situation represented by the analytical model. We conclude that,

in our scenarios, energy can be saved if the participants download more than ten files per week.

The future work will include the further development of shared applications and sharing patterns for our simulation environment. We aim at identifying additional thresholds that help to distinguish useful sharing from sharing that actually consumes more energy than it saves while considering the introduced constraints.

We furthermore will extend and implement the VHE economical fairness model, which is able to balance the usage of different resources like access network bandwidth, CPU usage, or sensor sharing [3].

**Acknowledgements** This project was partly funded by the German Research Foundation (Deutsche Forschungsgemeinschaft—DFG), contract number ME 1703/4-2 and by the EURO-FGI/Euro-NF—Network of Excellence, European Commission grant 028022/216366.

## References

1. Anderson T, Roscoe, T (2006) Learning from PlanetLab. In: Proceedings of the 3rd conference on USENIX workshop on real, large distributed systems
2. Bolch G, Greiner S, De Meer H, Trivedi K (2006) Queueing networks and markov Chains, 2nd edn. Wiley
3. Garcia AE, Berl A, Hummel KA, Weidlich R, Houyou AM, Hackbarth KD, De Meer H, Hlavacs H (2008) An economical cost model for fair resource sharing in virtual home environments (to be published in proceedings EuroNGI conference)
4. Hlavacs H, Hummel KA, Weidlich R, Houyou AM, De Meer H (2007) Modeling energy efficiency in distributed home environments. Technical report MIP-0714, Fakultät für Informatik und Mathematik, Universität Passau, December
5. IBM (2007) IBM virtualization view: virtualization can help power efficiency (2007) URL <http://www-03.ibm.com/systems/virtualization/view/011607.html>, January
6. Intel (2007) Energy Star\* System Implementation URL [www.intel.com/cd/channel/reseller/asm-na/eng/339085.htm](http://www.intel.com/cd/channel/reseller/asm-na/eng/339085.htm)
7. Intel (2006) Increasing data center density while driving down power and cooling costs URL [http://www.intel.com/business/bss/infrastructure/enterprise/power\\_thermal.pdf](http://www.intel.com/business/bss/infrastructure/enterprise/power_thermal.pdf) June
8. Intel (2006) Intel white paper 30057701: wireless intel SpeedStep power manager: optimizing power consumption for the intel PXA27x processor family (2004) URL <http://sunsite.rediris.es/pub/mirror/intel/pca/applicationsprocessors/whitepapers/30057701.pdf>
9. Koomey J (2006) Energystar, server energy measurement protocol, version 1.0, Following energy efficiency server benchmark technical workshop, Santa Clara, CA
10. Koomey J (2007) Estimating total power consumption by Servers in the US and the world. Tech. rep., Lawrence Berkeley National Laboratory and Stanford University

11. Peterson L, Roscoe T (2006) The design principles of PlanetLab. In: ACM SIGOPS operating systems review, vol 40, no 1. pp. 11–16
12. Rivoire S, Shah M, Ranganathan, P, Kozyrakis, C (2007) JouleSort: a balanced energy-efficiency benchmark. In: Proceedings of the 2007 ACM SIGMOD international conference on management of data
13. Stokes J (2007) Power plays: how power consumption will shape the future of computing. Online article: ars technica. URL <http://arstechnica.com/articles/paedia/cpu/power-plays.ars/2>
14. Thain D, Tannenbaum T, Livny M (2002) Condor and the grid. In: Grid computing: making the global infrastructure a reality. Wiley (March)
15. Weiser M (1999) The computer for the 21st Century. In: ACM SIGMOBILE mobile computing and communications review: special issue dedicated to Mark Weiser, vol 3, no. 3, pp 3–11, July
16. Windeck, C (2007) Energy star 4.0. C't German magazine for computer techniques, vol 14, pp 52–53
17. Windeck C (2007) Spar-o-matic. C't German magazine for computer techniques, vol 15, pp 200–207