Energy-Efficient Management of Physical and Virtual Resources - A Holistic Approach

Gergö Lovász, Florian Niedermeier, Hermann de Meer
University of Passau
Innstrasse 43, 94032 Passau, Germany
{gergoe.lovasz,florian.niedermeier,demeer}@uni-passau.de

***Abstract*—The spreading of information and communication technology has contributed much to the reduction of energy consumption in many areas of everyday life. Nevertheless the energy consumption of information and communication technology itself is rapidly growing and has to be dealt with. Currently used approaches focus mostly on the reduction of hardware energy consumption. This paper presents a vision of a holistic approach for reducing energy consumption in future communication infrastructures. Beside energy-efficient hardware as well as protocols that support the energy-efficient operation of communicating devices, the main focus of this paper is energy-efficient resource management. According to the Principle of Economic Efficiency, a limitation of resource provision is suggested by encapsulating applications in virtual machines with fixed resource requirements, together with the determination of an energy-minimal subset of resources on which applications are consolidated and which is able to fulfill the application requirements without over-provisioning.**

I. Introduction

The energy consumption of IT and communication infrastructures is dramatically increasing. A doubling of energy consumption from 2000 to 2005 of volume, mid-range, and high-end servers in the U.S. and worldwide is reported by Koomey [1]. In the year 2000 the total energy consumption of servers, routers, and PCs in Germany was about 5 billion kWh. For the year 2010 an energy consumption of more than 55 billion kWh per year is expected for the information and communication technology (ICT) in Germany [2].

Nowadays there are several approaches that address energy-efficient computing and communication. The most common approach is to develop more efficient hardware that consumes less energy and offers special energy efficiency features, e.g. energy saving modes. This effort is fostered by labels like the US Energy Star[1] or the European TCO Certification[2] who rate IT products (mostly monitors) for their environmental properties. But the development of energy-efficient hardware alone could not slow down or even stop the trend of increasing ICT energy consumption. Whereas it is important to develop hardware with low energy consumption and energy efficiency features, a holistic approach is needed to make ICT ”green”. A new energy-aware resource management has to be applied which considers the quality of service (QoS) requirements of applications on the one hand but also takes into account the energy consumption that is needed to perform a certain task on the other hand. Such a resource management aims at a resource allocation that is able to provide the QoS that is required by the applications and at the same time minimizes the energy that is needed to provide the requested services. Also energy-efficient application layer and communication layer protocols are needed, that support the energy-efficient operation of the hardware and an energy-efficient resource management by taking into account the energy-saving features of communicating devices, like power-saving modes. Section 2 gives a brief overview on energy-efficiency features of hardware, section 3 describes the problems with the currently used resource management paradigm and presents an alternative, energy-efficient way of managing resources. Section 4 describes the vision of an autonomous and energy-efficient resource management. Section 5 points out the weaknesses of currently used communication protocols with respect to energy efficiency and suggests methods to enable an energy-efficient communication. A conclusion is given in section 6.

[1] http://www.energystar.gov
[2] http://www.tcodevelopment.com

II. Energy-Efficient Hardware and its Features

Computer power can be saved by means of various techniques. First, the processor can be powered down by mechanisms like SpeedStep [3], PowerNow[3], Cool'n'Quiet[4], or Demand-Based Switching[5]. These measures enable slowing down the clock speeds (Clock Gating), or powering off parts of the chips (Power Gating), if they are idle [4],[5]. By sensing user-machine interaction, different hardware parts can incrementally be turned off or put in hibernating mode (display, disk, etc.). The ACPI specification defines four different power states which an ACPI-compliant computer system can be in. These states reach from G0-Working to G3-Mechanical-Off. The states G1 and G2 are subdivided into further states that describe which components are switched off in the particular state. For devices and the CPU separate power states (D0-D3 for devices and C0-C3 for CPUs) are defined which are similar to the global power states[6]. Some of the mentioned techniques are usually applied to mobile devices but can be

[3] http://www.amd.com/us-en/Processors/ProductInformation/0,,30_118_10220_10221%5E964,00.html
[4] http://www.amd.com/us-en/Processors/ProductInformation/0,,30_118_9485_9487%5E10272,00.html
[5] http://softwarecommunity.intel.com/articles/eng/1611.htm
[6] http://www.acpi.info/DOWNLOADS/ACPIspec30.pdf

used for desktop PCs as well. Energy saving techniques are also adopted by data centers to reduce power consumption. The overall power consumption of a data center consists of two major components: the power consumption of routers, servers, storage, switches and other devices on the one hand and the power consumption of the air-conditioning system (for hardware cooling) on the other hand. The energy efficiency standards of the hardware which is used in data centers exceeds the energy efficiency of ordinary PCs and hardware for personal usage in most cases.

III. The Principle of Economic Efficiency - A Resource Management Paradigm

The Principle of Economic Efficiency states two paradigms: The maximization and the minimization principle. The first aims at maximizing profit by optimizing the amount that is being produced with a given set of resources. The latter works in the opposite direction, minimizing the resources needed to provide a predefined output [6]. By applying the minimization principle on IT infrastructure, a massive gain in energy-efficiency can be expected.

A. Resource management using the maximization principle

Current resource management relies on the paradigm of maximization. This means that at any time, managed resources will provide a maximum of performance whereas energy (resource) consumption is in most cases only a minor concern. The hardware will always deliver the maximum available performance in order to provide the best service possible and scales its input power accordingly. Traditionally, application execution speed is only limited by the hardware capabilities. Application developers therefore allocate resources in a conservative way, meaning over-provisioning of resources in most cases and not taking into account the real requirements of the applications.

B. Resource management using the minimization principle

By performing the transition to the minimization principle, a fundamentally new paradigm is created. Managed resources should not provide maximum performance at any given time, but instead only deliver the amount of performance that is really needed by using as few resources as possible. Two key steps are necessary to implement the minimization principle:

1) Predefinition of the output
Often, especially in times with no or little load, the performance of components can be reduced without affecting user experience (e.g. lower connection speed, lower CPU frequency). In load-balanced environments, even complete machines may be shut down. By deploying services inside a VM, it is possible to effectively and predictably predefine the resource impact of a service. This predictability enables a resource management system to achive high utilizations in data center machines by deploying VMs in a "‘best fit"’ way. By carefully examining the required level of QoS, VM parameters can be set according to real resource needs.
2) Minimization of the input
Finally, knowledge about the resource requirements of the applications and capabilities of the underlying hardware are used to find a near energy-minimal subset of resources. This subset is defined as a set of hardware resources that provides all requested services (within the predefined QoS requirements) at a given time, while consuming the least possible amount of energy in a defined environment.
Let R be the set of available hardware resources in a certain environment (e.g., a federation of data centers, or a cloud computing environment) and let A be the set of applications (e.g., a mail or web service) with predefined QoS requirements that have to be provided. Let $energy(A, X)$, with $X \subseteq R$ be the energy needed to run A on a subset of hardware resources X. $E \subseteq R$ is an energy-minimal subset of hardware resources for A if
    a) E can provide the required QoS to A,
    b) $energy(A, E) \leq energy(A, F)$, where F, $(F \subseteq R)$ is any other subset of hardware resources that can provide the required QoS to A.

If two subsets of hardware provide the requested services in an energy-minimal way, the preferred subset is the set that provides the best QoS. Finding $E \subseteq R$ is an instance of the multidimensional bin-packing problem and therefore an NP-hard problem. However, even an approximation of $E \subseteq R$ could lead to a considerable reduction of energy consumption since resources that are not needed can be turned off or set into an energy-saving mode. Recently, a feature called live migration is becoming available in an increasing number of VMMs. This feature enables a nearly seamless migration of a running service inside a virtual machine from one physical host to another. Therefore it is possible to dynamically determine the energy-minimal subset of resources for a given set of applications and reallocate resources by consolidation of applications at runtime.

IV. Energy-efficient Resource Management

In this section a vision of an autonomic and energy-efficient monitoring is presented which realizes a resource management according to the minimization principle and is based on energy-related monitoring and energy consumption models. Figure 1 shows the main components of such a resource management system.

Main components are energy consumption models of the managed devices (e.g. servers, routers) that are needed to compute estimations of the energy consumption of different resource allocations. A model describes the energy consumption of a resource depending on the load of its sub-components. Therefore the model is devided into static properties and dynamic properties. Whereas the static properties describe the characteristics of hardware and usually do not change (e.g.

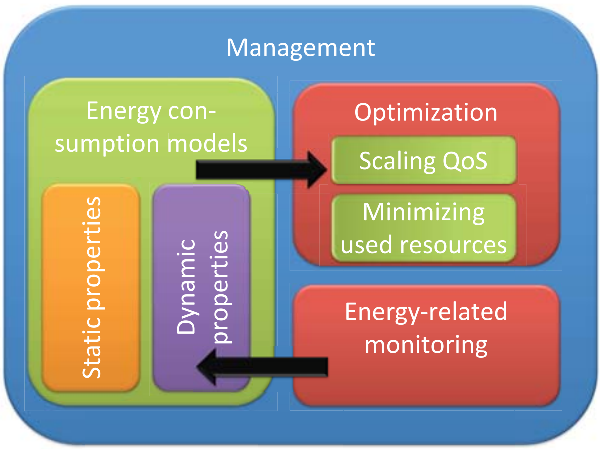

Fig. 1. Energy-efficient management of resources based on energy consumption models and energy-related monitoring

number of CPU cores, size of RAM), the dynamic properties strongly depend on the load on the resource and its components (e.g. the load on a server's CPU, used RAM). To identify the energy-relevant components and properties of a hardware resource, detailed measurements have to be performed. In a first stage, the single components are stress tested and the influence on the total energy consumption is measured. In a second stage, the cross-correlation of different components is analyzed. By combining the results of both stages, a generic energy consumption model can be derived for the resource that will have static and dynamic input parameters.

Another main component of the resource management system is the energy-related monitoring of resources. Realtime monitoring has to be applied to provide the dynamic, energy-relevant monitoring data as input to the energy consumption models. This includes especially the load of certain resource components (e.g. CPU, HDD). But also environmental properties, like the temperature, have to be monitored. This information can then be used for the distribution of server load to avoid hot spots in the equipment and to minimize the energy consumption of the air conditioning.

A key component of the management system is the optimizer. Through the encapsulation of applications in virtual machines it is able to exactly determine the application requirements (predefined output). The energy consumption models and the energy-related monitoring allows the optimizer to compute different potential resource allocations, estimate their power consumption, and approximate the energy-minimal subset of resources (minimal input).

V. Energy-Efficient Communication Protocols

Another critical area that has to be addressed in order to reduce ICT's energy consumption are communication protocols. Whereas communicating devices often offer mechanisms to save energy, communication protocols prevent the application of certain energy-saving features. Communication protocols have to be analyzed, concerning the energy-efficient support of network elements and their features. The protocols have to be designed in a way that the availability of certain network services does not assume the permanent accessibility of all virtual or physical entities. This enables devices to change into an energy-efficient mode for certain time periods. Protocols have to support the synchronization of communication, the delegation of services, on-demand mechanisms, and an energy-efficient signaling. Not all services do necessarily require permanent and direct accessibility, for instance.

- **Synchronization of communication**
  Communication between two entities can be synchronized in such a case, to allow the devices to hibernate between two communication phases without interrupting the availability of the service. Protocol mechanisms have to be developed which allow the synchronization of active communication phases. Devices should hibernate not only when no communication is taking place but also if they are lowly utilized (e.g. persistent TPC connections).
- **Delegation of services and functions**
  A network device should be enabled to delegate services to other devices, if possible (clients should not be aware of this delegation). Such a delegation would allow the transfer of a service from an energy inefficient to a rather energy efficient device or to a device which has to be always on anyway (e.g. a router). The delegating device can change to dormant mode or can be turned off. The delegated device can either provide the services itself or wake up the delegating device in case of a request. In order to keep the response time low for the requesting device in the latter case, the delegated device can take over first communication steps (e.g. connection establishment) while waiting for the delegating device to change to active mode. Alternatively, nodes can be activated on demand, instead of being permanently available.
- **Get rid of heart-beats, power on network devices on demand**
  Current protocols, like TCP or BGP, rely on keep-alive messages to determine the reachability of a host. These keep-alive messages prevent hardware from entering sleep states while at the same time being available in the network. The protocol stack has to be modified to not assume constant reachability of each device. Instead, device need to be able to enter a new state that allows them to be marked as available, but on standby. This state should allow machines to enter a sleep mode and only be waked up if actual communication with this machine is attempted. Mechanisms have to be developed that allow entities in the network to buffer context information in order to minimize signaling when a sleeping device changes to active mode.
- **Signaling**
  Currently, signaling and data traffic share the same physical links, although they have significantly different key features. While data traffic is likely to need high bandwidth and occur in bursts, signaling packets are very

small in size and more evenly spread over time. Using the same link for both means that in idle periods the capacity of the link is much larger than needed for signaling. Using a seperate, low bandwidth link exclusively for signaling would allow the high performance link to be shut down in idle periods to save energy. Alernatively, variable speed connections could be established between the network devices, allowing a low speed to be negotiated in idle periods. This lower speed of the link could in turn enable processors used inside switches and computer NICs to scale down their frequency.

VI. Conclusion

To counter the rapid increase in ICT energy consumption, a new paradigm has to be introduced. The transition from maximization to minimization principle is driven by the restriction of output to the needed levels and minimization of the power consumption. We implement this principle by using virtual machines to encapsulate services and consolidate them onto a near energy-minimal subset of the given physical resources. To achive an autonomous management of the infrastructure, energy consumption has to be modeled and predicted using real-time monitoring data. To achieve a further reduction of energy consumption, current communication protocols have to be replaced by energy-aware protocols to allow devices to use sleep states more effectively.
Future research will include a prototype implementation of the envisioned management system on the G-Lab testbed.

Acknowledgment

The authors would like to thank The German Ministry for Education and Research (BMBF) who is funding the project G-Lab_Ener-G that deals with energy efficiency in future networks.

References

[1] Koomey. Estimating total power consumption by servers in the us and the world, Technical report, 2007, http://enterprise.amd.com/Downloads/svrpwrusecompletefinal.pdf

[2] Fraunhofer ISI. Der Einuss moderner Gertegenerationen der IuK-Technik auf den Energieverbrauch in Deutschland bis zum Jahr 2010, Studie. http://www.isi.fhg.de/e/publikation/iuk/Fraunhofer-IuK-Kurzfassung.pdf

[3] Intel white paper 30057701. Wireless Intel SpeedStep Power Manager: Optimizing Power Consumption for the Intel PXA27x Processor Family, sunsite.rediris.es/pub/mirror/intel/pca/applicationsprocessors/whitepapers/30057701.pdf, 2004

[4] ENERGY STAR* System Implementation, Published by Intel with technical collaboration from the U.S. Environmental Protection Agency, Whitepaper, 2007, Revision-001, http://www.energystar.gov/ia/partners/prod_development/revisions/downloads/316478-001.pdf

[5] Windeck. Energy Star 4.0, C't German Magazine for Computer Techniques, Vol. 14, 52-53, 2007

[6] E. Grochla and E. Gaugler, Eds., Handbook of German Business Management. Poeschel, Stuttgart, Springer, Berlin/Heidelberg, 1990, vol. 1 (A-E).