European Transactions on Telecommunications

Power Efficiency of Thin Clients

Journal:	*European Transactions on Telecommunications*
Manuscript ID:	ETT-08-0118
Wiley - Manuscript type:	Research Article
Date Submitted by the Author:	25-Sep-2008
Complete List of Authors:	Vereecken, Willem; UGent - IBBT - IMEC, INTEC - IBCN Deboosere, Lien; UGent - IBBT - IMEC, INTEC - IBCN Simoens, Pieter; UGent - IBBT - IMEC, INTEC - IBCN Vermeulen, Brecht; UGent - IBBT - IMEC, INTEC - IBCN Colle, Didier; UGent - IBBT - IMEC, INTEC - IBCN Develder, Chris; UGent - IBBT - IMEC, INTEC - IBCN Pickavet, Mario; UGent - IBBT - IMEC, INTEC - IBCN Dhoedt, Bart; UGent - IBBT - IMEC, INTEC - IBCN Demeester, Piet; UGent - IBBT - IMEC, INTEC - IBCN
Keywords:	Thin Client, Energy Efficiency, Power Consumption, New Network Paradigms

EUROPEAN TRANSACTIONS ON TELECOMMUNICATIONS
Euro. Trans. Telecomms. **00**: 1–11 (0000)

(www.interscience.wiley.com) DOI: 10.1002/ett.0000

Power Efficiency of Thin Clients

Willem Vereecken[1] Lien Deboosere[1] Pieter Simoens[1] Brecht Vermeulen[1]
Didier Colle[1] Chris Develder[1] Mario Pickavet[1] Bart Dhoedt[1] Piet Demeester[1]

[1]*Ghent University - IBBT, Department of Information Technology (INTEC)*
Gaston Crommenlaan 8, Bus 201, 9050 Ghent, Belgium
Email: {firstname.lastname}@intec.ugent.be

SUMMARY

Worldwide, awareness for energy consumption is raising because of global energy production limits as well as because of environmental concerns. As the energy fraction currently consumed by ICT related equipment is substantial (about 8% of electricity production worldwide in the use phase) and the growth rate in this particular sector is spectacular, also in the ICT sector adequate solutions are needed to allow sustainable growth. This paper analyzes the conditions for thin client computing to significantly reduce energy consumption. To this end, estimates on power consumptions in typical desktop scenarios and analogous thin client settings are made and analyzed. The paper concludes with an experimental study on currently available equipment, to translate the generic conclusions into their current implications and trade-offs. Copyright © 0000 AEIT

1. Introduction

Only recently awareness is raising that power consumption directly related to ICT equipment and services represents a relevant fraction of the worldwide energy production. Given the annual growth rate of these services, in some cases exceeding 10% on a yearly basis, ICT related power consumption is indeed becoming an increasingly worrying concern. As more and more businesses are relying on sustained ICT services, energy concerns might constrain economical growth in a number of vital economical sectors, thereby jeopardizing the wealth to a considerable extent.

The explosive growth in ICT related energy consumption can be explained by a number of trends: not only the worldwide adoption rate of existing services (including broadband Internet services, mobile communication services, ...) is to blame, but also the emergence of new, resource and energy hungry services. Amongst the latter category, an important example is the birth of upload and consumption services of personal content (still images and video), requiring huge data centers and high speed network facilities. Also replacement of existing equipment by state-of-the-art devices generally implies a substantial increase in power consumption: over the last years the average home PC power consumption has increased by more than 10% annually. [1] [2]

The combination of these mechanisms has brought us in a situation where the ICT related energy consumption can be estimated at 4% of the primary energy production in 2008. Forecasts for 2020 are typically in the range of 8% [3]. It is clear that this growth rate will be difficult to sustain, also in view of rising energy prices in combination with environmental concerns.

Solutions to save power almost always use the same underlying technique: scale down the performance of devices as much as possible or even shut down equipment when possible. This technique is well known in mobile computing, arguably the sector where terminal power consumption is of prime importance (to improve battery lifetime): wireless transmission protocols switch to a lower transmission speed when possible (and go to standby mode), and the wireless card is even shut down when no network activity is detected. A second domain where power consumption is of prime importance, is the data

Prepared using ettauth.cls *[Version: 2007/01/05 v1.00]*
Received September 25, 2008

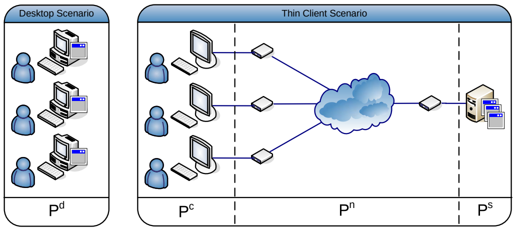

Figure 1. Desktop and Thin Client Scenario

center application area. In this area it is worth mentioning the Green Grid consortium[4] focussing on advancing energy efficiency in data centers.

Despite of the critical importance of the problem, relatively few initiatives are underway to identify adequate solutions guaranteeing sustainable growth in the ICT sector. On the network level, a good example is the IEEE study group on Energy Efficient Ethernet [5], where power savings for Ethernet are studied, again based on scaling down the link bit rate in a coordinated way. Another important avenue for power saving is adopting the thin client computing paradigm [6]. This approach is in fact not unlike the mainframe approach generally adopted in the '60s-'70s (and left again in the early '80s), where a server farm is performing the computational intensive (and hence energy hungry) functions, while the rendering for the end-user is done on very constrained devices.

In this paper we will demonstrate the thin clients paradigm can increase power-efficiency. In sections 2 & 3 we will analyse the active state and passive state power consumption of the thin client paradigm. In section 4 we will evaluate the model based on real world measurements. In section 5 the major conclusions will be summarized.

2. Active state analysis

When evaluating the power efficiency of the thin client paradigm we consider two scenarios. On the one hand we have a traditional desktop where each user is running a standalone application on a standard PC. In the second scenario the desktops are replaced with thin client terminals. The standalone applications run remotely on servers in the data center. Both scenario's are schematically depicted in Fig. 1.

It is already clear that the thin client scenario has a number of advantages and disadvantages when compared to the desktop scenario:

+ Power consumed by the thin client terminal is significantly lower than a normal desktop PC.

+ Server side resources can be delivered more efficiently: high-end servers are shared between all users, implying that the server infrastructure will exhibit less idle periods (the typical load on desktops is below 20%).

– Applications are run remotely, implying possibly application specific network overhead (e.g. for sending input events to the server and getting screen updates back). Additional equipment is needed (e.g. switch at server side, network interface cards consuming power ...).

– Protocol overhead from the thin client protocol requires additional server side processing.

– Resources at the server side must be cooled, increasing the power budget for the thin client scenario. Given these observations, one can already conclude that the balance for the thin client paradigm will certainly depend on the following factors: server resource efficiency (influenced by the achievable amount of sharing and optimal resource usage), server side cooling efficiency and the bandwidth consumption (assuming power scales with consumed bandwidth, it is clear that high bandwidth applications, such as e.g. multimedia editing will benefit less).

2.1. *Desktop Scenario*

When in the active state, the main sources of power consumption are the CPU, the hard disk and the network interface. (Note that power consumption caused by the monitor is not taken into account, as a similar amount of power would be consumed in the thin client monitor.) Each of these hardware elements is characterized by a load (a real number between 0 and 1), i.e. λ^d_{CPU}, λ^d_{HD} and λ^d_{NIC} denoting the load on the CPU, the hard disk and the network interface card respectively. The unloaded power consumption for each of these is written as $P^d_{0,CPU}$, $P^d_{0,HD}$ and $P^d_{0,NIC}$ and the power consumed in loaded conditions is therefore (with '*' representing CPU, HD or NIC):

$$P^d_* = P^d_{0,*} + f^d_*\left(\lambda^d_*\right) \tag{1}$$

Note that the CPU power consumption incurred by network traffic is included in P^d_{NIC} . The functions f^d_* simply express the relation between device load and power consumption.

The overall power consumption of single desktop therefore equals

$$\begin{aligned} P^d &= \sum_{*=CPU,HD,NIC} \left[P^d_{0,*} + f^d_*\left(\lambda^d_*\right)\right] && (2)\\ &= P^d_0 + \sum_{*=CPU,HD,NIC} f^d_*\left(\lambda^d_*\right) && (3)\end{aligned}$$

DOI: 10.1002/ett

with P_0^d the total unloaded power consumption of the desktop:

$$P_0^d = P_{0,CPU}^d + P_{0,HD}^d + P_{0,NIC}^d \tag{4}$$

Experimental measurements show the power consumed by the hard disk is heavily dominated by the rotation motor of the drive, and far less by the load. In addition the relation f_{CPU}^d is linear. Therefore we assume a factor of α_{CPU}^d. Since we assume a standalone application, the network card is unloaded. The desktop power consumption becomes:

$$P^d = P_0^d + \alpha_{CPU}^d \lambda_{CPU}^d \tag{5}$$

2.2. Thin Client Scenario

In the thin client several types of equipment need to be considered. Firstly we will consider the client terminal and the server. These will behave like a desktop PC. However we also need to consider the power consumption caused by the load λ_{NIC}^* on the network interface card (NIC). Secondly the power consumption of the network needs to be modelled. Thirdly we need to take into account that certain equipment types are located in a data center. This equipment is being cooled and the power consumption of the cooling also needs to be incorporated in the model.

2.2.1. The Client Terminal A thin client terminal typically behaves like a desktop PC without a hard drive. After experimental measurements the power consumption appears to be constant even with varying CPU load λ_{CPU}^c and NIC load λ_{NIC}^c. Therefore the power consumption is reduced to:

$$P^c = P_0^c \tag{6}$$

2.2.2. The server Again based on experimental measurements we assume linear dependencies in the power consumption for the server. The power consumption becomes:

$$P^s = P_0^s + \alpha_{CPU}^s \lambda_{CPU}^s + \alpha_{NIC}^s \lambda_{NIC}^s \tag{7}$$

The load λ_{NIC}^s is in reality the bandwidth received by the server b^s. We express this bandwidth as a function of the bandwidth b percieved at the client. When assuming a share ratio of N users per server we get:

$$\lambda_{NIC}^s = b^s = Nb \tag{8}$$

Obviously, the load on the server λ_{CPU}^s is related to the load on the clients. The amount of work to be performed by a single server is at least the amount of work done by N desktops. On the other hand, there is the extra work needed on the server to support N sessions, and processing the protocol overhead (to receive input from the thin clients and to construct and send back screen updates). If we note the processing capacity of a server (according to a relevant performance oriented benchmark such as SPECint2000 [7]) as C^s and the analogous parameter for the desktop case C^d we have

$$\lambda_{CPU}^s C^s > N \lambda_{CPU}^d C^d \tag{9}$$

By denoting the extra load caused per user by ϵ, we have

$$\lambda_{CPU}^s = N\left[\lambda_{CPU}^d \frac{C^d}{C^s} + \epsilon\right] \tag{10}$$

Since $\lambda_{CPU}^s \leq 1$ we get a maximal value for N:

$$N \leq \left[\lambda_{CPU}^d \frac{C^d}{C^s} + \epsilon\right]^{-1} \tag{11}$$

2.2.3. The network There are several possibilities to model the network power consumption. In this study we have chosen to make abstraction of the specific layout of the network. We denote the network power consumption per user as P^n. P^n represents the fully allocated power consumption of the following devices:

- The LAN switch in the client network
- The gateway in the client network
- The wide area network devices (routers and traffic aggregators)
- The gateway in the data center
- The LAN switches in the data center

P^n is determined by the created traffic load on the network. This load is the bandwith b caused by the thin client protocol. Based on measurements performed on several network devices we assume again linear relations and get:

$$P^n = P_0^n + \alpha_T^n b \tag{12}$$

2.2.4. Cooling Due to the concentration of heat dissipating equipment, considerable efforts are needed to cool data centers. This cooling infrastructure of course also consumes electrical power. Therefore not all electrical power consumed by the datacenter is used for the ICT equipment. This factor is denoted by the Power Usage Effectiveness (PUE) [8]:

$$PUE = \frac{P^{dc}_{tot}}{P^{dc}_{ICT}} \qquad (13)$$

Since our model should cover multiple cases we will consider the PUE accounted for in the relevant parameters.

2.2.5. Total The total power of the setup, assuming N^u users is:

$$N^u P^c + N^u P^n + \frac{N^u}{N} P^s \qquad (14)$$

In order to compare the power consumption with the desktop scenario we need to divide this by N^u. The power consumed for one thin client is:

$$P^{tc} = P^c + P^n + \frac{1}{N} P^s \qquad (15)$$

Substitution in this formula leads to:

$$P^{tc} = P^c_0 + P^n_0 + (P^s_0 + \alpha^s_{CPU}\lambda^s_{CPU})\frac{1}{N} + (\alpha^n_T + \alpha^s_{NIC})\, b \qquad (16)$$

3. Passive State analysis

In the previous section, it was assumed that all clients were in the active state. In this section we will study potential benefits and drawbacks arising from passive clients. Two mechanisms contribute to reduced power consumption in the client scenario:

- The thin client terminal consumes less power when off-line reducing power consumption in the passive state at the client side.
- Servers can be put in a sleep mode when a number of users go to the passive state, thereby reducing the power consumption in the data center.

The total number of users is still denoted by N^u. The quantities N^u_{act}, N^u_{off} and N^u_{sb} denote the number of clients in the active state, the off state and the standby state respectively. Power consumed by device '*' in these states is represented by P^*_{act}, P^*_{off} and P^*_{sb} respectively. Clearly, we have at any given moment (N^u constant)

$$N^u = N^u_{act} + N^u_{off} + N^u_{sb} \qquad (17)$$

as well as the average (averaging per user) power consumption of device '*' ($* = d, c$)

$$P^*_{avg} = \frac{N^u_{act}}{N^u} P^*_{act} + \frac{N^u_{off}}{N^u} P^*_{off} + \frac{N^u_{sb}}{N^u} P^*_{sb} \qquad (18)$$

Applying the approximations used in the previous sections for the active state power consumption, we find for the desktop and the thin client terminal:

$$P^d_{avg} = P^d_{0,avg} + \frac{N^u_{act}}{N^u}\alpha^d_{CPU}\lambda^d_{CPU} \qquad (19)$$

$$P^c_{avg} = P^c_{0,avg} \qquad (20)$$

Similary we get for the server

$$P^s_{avg} = \frac{N^s_{act}}{N^s} P^s_{act} + \frac{N^s_{off}}{N^s} P^s_{off} + \frac{N^s_{sb}}{N^s} P^s_{sb} \qquad (21)$$

If we denote the server CPU load in active state as $\lambda^s_{CPU,AS}$ given by (10) and the share ratio $N_{act} \triangleq N^u_{act}/N^s_{act}$ we get:

$$\lambda^s_{CPU} = N_{act}\left[\lambda^d_{CPU}\frac{C^d}{C^s} + \epsilon\right] \qquad (22)$$

$$\Rightarrow \lambda^s_{CPU} = \frac{N_{act}}{N}\lambda^s_{CPU,AS} \qquad (23)$$

This means:

$$P^s_{avg} = P^s_{0,avg} + \frac{N^u_{act}}{N^s}\frac{1}{N}\alpha^s_{CPU}\lambda^s_{CPU,AS} + \alpha^s_{NIC}\frac{N^u_{act}}{N^s} b \qquad (24)$$

In the network we also consider three states. In the active state the equipment is performing its full functionality. In the off state the equipment is switched of. In the standby state the equipment has a reduced functionality. Typically network devices in standby operate at reduced power with a bitrate of $128 kBit/s$. Additionally they have a small wake up time so the network functionality is not compromised.

At the user premises one can afford to switch off the network equipment. Deeper in the network this is however not possible. The user is not present to activate the equipment and moreover the equipment is shared between multiple users. Therefore defining the network state as switched off or standby is not as straightforward as with the desktops, client terminals and servers. We define a number of reduced power states for the network. In order to maintain the generality of the model we do not further define what the reduced power states entail. These reduced power states will only affect P^n_0 since the bandwidth b is only originating from the active connections. Thus P^n in the passive state becomes:

$$P^n_{avg} = \frac{N^u_{act}}{N^u} P^n_0 + \frac{N^u_{sb}}{N^u}\sum_i f^{n,sb}_i P^{n,sb}_{red,i} + \frac{N^u_{off}}{N^u}\sum_i f^{n,off}_i P^{n,off}_{red,i} + \alpha^n_T b \qquad (25)$$

Where $f_*^{n,*}$ denotes the fraction of the representation of a certain reduced power state $P_{red,*}^{n,*}$ for the client terminal in standby or switched off.

3.1. Desktop Scenario

When comparing the power consumption in the passive state 19 to the active state 5 it is obvious that the power consumption of the desktops is reduced by the power saving of the machines that are shut down or in standby and the CPU load not consumed by these machines.

3.2. Thin Client Scenario

For the thin client scenario we have yet to define the state distribution N_*^s of the servers and the passive state of the network. Firstly we will consider three scenarios with a fully active network (so no reduced power states). In the first scenario unused servers are not put in standby mode (or even shut down). In the second scenario we will assume that servers can be put in a power saving mode. In the third scenario we will shut down the servers instead of putting them in power saving mode. It is obvious that the second and the third scenario will imply power savings in the model. There are however some drawbacks to these scenarios:

- Reducing the number of active servers while sessions are running requires a flexible migration of these sessions in order not to affect the active users.
- Shutting down servers is less flexible than putting them in standby. The responsiveness of the server management under varying activity of the users will have to be evaluated.

Secondly we will introduce the additional power saving we can get from reduced power states in the network.

3.2.1. Scenario I: All servers remain active When all servers remain active $N^s = N_{act}^s$. This means:

$$N_{act}^s = N^s; \quad N_{sb}^s = 0; \quad N_{off}^s = 0 \tag{26}$$

$$N_{act} \triangleq \frac{N_{act}^u}{N_{act}^s} = \frac{N_{act}^u}{N^u} N \tag{27}$$

Using these values we can calculate the power consumption in this scenario:

$$\begin{aligned} P_I^{tc} = P^{tc} - & \\ & \left(\frac{N_{off}^u}{N^u} \left(P_{act}^c - P_{off}^c \right) + \frac{N_{sb}^u}{N^u} \left(P_{act}^c - P_{sb}^c \right) \right) \\ - \frac{N_{sb}^u + N_{off}^u}{N^u} & \left(\alpha_{CPU}^s \lambda_{CPU,AS}^s \frac{1}{N} + \left(\alpha_T^n + \alpha_{NIC}^s \right) b \right) \end{aligned} \tag{28}$$

P^{tc} is (16) with P_0^c and P_0^s given by $P_{0,act}^c$ and $P_{0,act}^s$ respectively.

The power consumption is reduced by two factors. Firstly we see the obvious power saving caused by the clients being shut down or in standby. Secondly the power consumption is further reduced by the lower CPU load and bandwidth consumption.

3.2.2. Scenario II: Servers in power saving mode when possible In this scenario, unused servers are put in a power saving. Since only a fraction of $\frac{N_{act}^u}{N^u}$ of users consumes computing cycles, we assume that only this fraction of servers is active. This translates into:

$$N_{act}^s = \frac{N_{act}^u}{N^u} N^s = \frac{N_{act}^u}{N}; \quad N_{sb}^s = N^s - \frac{N_{act}^u}{N}; \quad N_{off}^s = 0 \tag{29}$$

$$N_{act} \triangleq \frac{N_{act}^u}{N_{act}^s} = N \tag{30}$$

We get:

$$P_{II}^{tc} = P_I^{tc} - \frac{N_{sb}^u + N_{off}^u}{N^u} \frac{1}{N} \left(P_{0,act}^s - P_{0,sb}^s \right) \tag{31}$$

One sees that the power consumption is further reduced by the power saving of the servers in standby.

3.2.3. Scenario III: Servers shut down when possible In this scenario we use the same assumptions as in the previous section. Only now we shut down the servers instead of putting them in power saving mode. This translates into:

$$N_{act}^s = \frac{N_{act}^u}{N}; \quad N_{sb}^s = 0; \quad N_{off}^s = N^s - \frac{N_{act}^u}{N} \tag{32}$$

$$N_{act} \triangleq \frac{N_{act}^u}{N_{act}^s} = N \tag{33}$$

DOI: 10.1002/ett

We get:

$$P_{III}^{tc} = P_{II}^{tc} - \frac{N_{sb}^u + N_{off}^u}{N^u}\frac{1}{N}\left(P_{0,sb}^s - P_{0,off}^s\right) \quad (34)$$

A similar power saving as in the previous scenario can be found.

We can simplify these results by assuming that we physically cut off the power of shut down equipment. That means $P_{off}^* = 0$. Further we assume that the inactive thin clients and desktops are shut down so $N_{sb}^u = 0$. Note that these assumptions mean that we use both the desktop and the thin client solutions in their most energy efficient way.

$$\begin{aligned} P_{III}^{tc} = P^{tc} - \frac{N_{off}^u}{N^u}[P_{act}^c \\ + \left(P_{0,act}^s + \alpha_{CPU}^s \lambda_{CPU,AS}^s\right)\frac{1}{N} + \left(\alpha_T^n + \alpha_{NIC}^s\right) b] \end{aligned} \quad (35)$$

$$P_{III}^{tc} = \frac{N_{act}^u}{N^u} P^{tc} + \frac{N_{off}^u}{N^u} P_0^n \quad (36)$$

The power consumption of the thin client solution scales with the number of active users except for the basic network power consumption.

3.2.4. Scenario IV: Reduced Power states in the network

To further scale down power consumption the only remaining option is introducing reduced power states in the network. For the passive network connections we assume one reduced power state P_{red}^n. Using (25) we get:

$$P_{avg}^n = \frac{N_{act}^u}{N^u} P_0^n + \frac{N_{off}^u}{N^u} P_{red}^n + \alpha_T^n b \quad (37)$$

This leads to

$$P_{IV}^{tc} = \frac{N_{act}^u}{N^u} P^{tc} + \frac{N_{off}^u}{N^u} P_{red}^n \quad (38)$$

$$= P_{III}^{tc} - \frac{N_{off}^u}{N^u}\left(P_0^n - P_{red}^n\right) \quad (39)$$

It is clear that in order to achieve maximal energy effciency the P_{red}^n needs to be minimal. When $P_{red}^n = 0$ the energy consumption of the thin client solution scales with the number of active users. Note however that this case is only theoretical since we need to maintain a minimal connectivity in the network.

Desktop PC		Laptop PC	
P_0^d	$82.6W$	P_0^d	$28.6W$
α_{CPU}^d	$13.9W$	α_{CPU}^d	$10W$
C^d	1401	C^d	1541
Client Terminal		**Server**	
P_0^c	$4W$	P_0^s	$217W$
		$P_{0,sb}^s$	$15.8W$
		α_{CPU}^s	$10.42W$
		α_{NIC}^s	$0.93\frac{mW}{Mb/s}$
		C^s	4×1435
		PUE	2

Table 1. Equipment Parameters

4. Experimental Results

When evaluating practical implementations it is important to gain insight in the power saved by implementing thin clients. Therefore we will evaluate two parameters. We define the saved power as $\Delta P = P^d - P^{tc}$ which will express the power saving for a single user. The second parameter is the power ratio $R = \frac{Pd}{Ptc}$ which expresses the relative power saving between both scenarios. The criterium for power efficiency is:

$$\Delta P > 0 \quad (40)$$

or stated otherwise:

$$R > 100\% \quad (41)$$

We evaluate these parameters in function of the server share ratio N and the network power consumption P_0^n. The average load λ_{CPU}^d will be approximately 20% which is largely sufficient for standard office applications such as text editors and spreadsheets. Since for these applications network connectivity is not required we consider no network power consumption in the desktop scenario. In the thin client scenario the standalone applications run remotely on servers in the data center. The consumed bandwidth will vary between $0Mb/s$ and $5Mb/s$[9]. The server overhead ϵ is considered to be small ($\epsilon \approx 0$). The servers are located in a data center for which we assume a typical PUE of 2.

We measured the power consumption of a desktop (AMD Athlon 64 3500+™), a laptop (Pentium M™ 2GHz), a server (AMD Opteron 2212™) and a thin client device (JackPC™). The measured parameters are summarized in table 1. The profile of the power consumption of the server corresponds with the typical

DOI: 10.1002/ett

	User Prem. Eq.	Access Netw. Eq.	Total
Active state			
ADSL2	1.5 W	1.2 W	2.7 W
VDSL2	6.0 W	1.6 W	7.6 W
PON	12.0 W	0.2 W	12.2 W
Reduced Power State			
ADSL2	0.0 W	0.8 W	0.8 W
VDSL2	0.3 W	1.0 W	1.3 W
PON	0.3 W	0.2 W	0.5 W

Table 2. Power Consumption per User of Network Equipment [11]

behaviour as can be seen in [10]. The bandwith factor α^s_{NIC} appears to be small compared to the relevant bandwith and the other parameters. We assume the same order of magnitude for α^n_T Therefore we will ignore the factor $(\alpha^n_T + \alpha^s_{NIC})\, b$.

In order to limit the network latency we assume the data center to be located in the access network. [11] mentions target values for the power consumption of the network equipment. We consider three network technologies: ADSL2, VDSL2 and PON. For the access network power consumption of the PON we assume a typical value of $0.2W/subscriber$. For the reduced power state we assume the equipment at the user premises to be switched off and the access network equipment in standby state. The network power consumption values are summarized in table 2.

The power consumption share of network equipment deeper in the network is not accounted for in this case. First of all because, as stated before the servers cannot be too deep in the network. Moreover, that equipment will be shared over a large number of users so the power consumption per subscriber becomes negligable. Note as well that this is why the PON case is the least power efficient. The bandwidth provided by this solution is significantly larger than for the other solutions whereas this is not required. It would be more efficient to share these high bandwidths over a larger number of users and then implement the final connection with an ADSL2/VDSL2 line. This case however is covered by the first two solutions.

We analyse the active state and the passive state. To limit the complexity we will assume that the desktops or client terminals are either active or switched off ($N^u_{sb} = 0$).

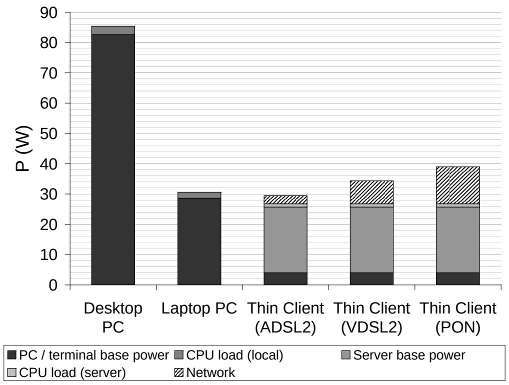

Figure 2. Power consumption of Desktop PC, Laptop PC and Thin Client in active state for $\lambda^d_{CPU} = 20\%$ and N=20

4.1. Active State Analysis

Fig. 2 displays a breakdown in the power consumption for a Desktop PC, a Laptop PC and a Thin Client Setup. Compared to the Desktop PC the power consumption of Thin Client Setup is significantly lower. When comparing with a Laptop PC the power consumption of both solutions is in the same order of magnitude for ADSL2. When using VDSL2 or PON the thin client solution is even disadvantageous. Note however that manufacturers limit the power consumption of a laptop as much as possible. When comparing P^d_0 one can see a laptop PC is roughly three times more efficient than a desktop PC. This for two machines that have the same functionality. Similar optimizations should be possible for the servers and to a lesser degree the network equipment. The basic power consumption of the server P^s_0 accounts for the largest amount of the power consumption of the thin client solution. The previously mentioned optimizations in server technology and reducing the PUE of the data center should enable the thin client paradigm to become significantly more power efficient than the laptop PC. To obtain a fair comparison between both scenarios we will focus on comparing the desktop PC with the thin client solution.

Fig. 3 displays ΔP and R in function of the server share ratio N for the different technologies. Next to the three network technologies the figures include the theoretical case of $P^n = 0$ as well. The power saving is highly dependent on the share ratio on the servers. A minimal share ratio of $N > 5$ in order to be more efficient than the desktop scenario. At the maximal share ratio of $N \approx 20$ ($\lambda^s_{CPU} \approx 1$) power savings up to $50W$ (300%) are possible. It is also clear that the choice of network technology can have an important impact on the power

Prepared using ettauth.cls

DOI: 10.1002/ett

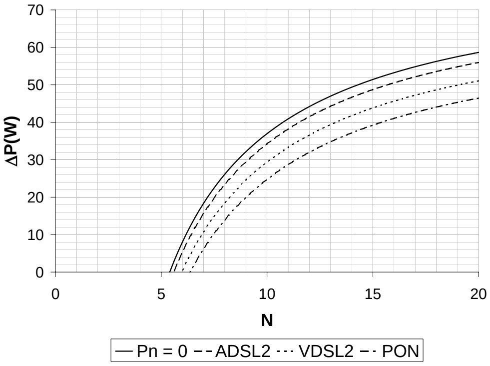

(a) ΔP

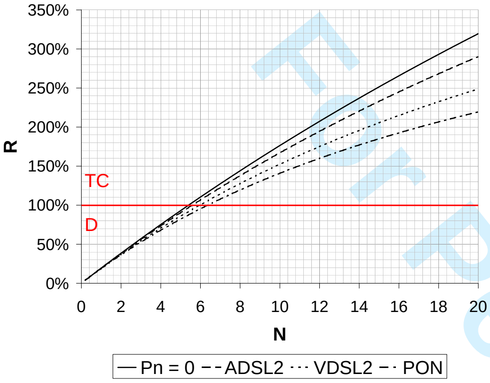

(b) R

Figure 3. Power saving of Thin Client towards Desktop PC in function of Server Share Ratio and Network Power Consumption

saving possibilities. However, the impact of the server power consumption stil remains the most significant.

4.2. *Passive State Analysis*

When regarding the passive state analysis the three relevant scenario's are displayed in Fig. 4. For λ^d_{CPU} we assumed a value of 20%. We assumed a share ratio of $N = 20$. However, the conclusions are qualitatively similar for $N < 20$.

When all the servers remain active (I) the efficiency degrades approximately linearly in function of the fraction of passive users $\frac{N^u_{off}}{N^u}$. Putting servers in standby (II) or switching them off (III) can lead to large increases in the efficiency. If the network power consumption is low (ADSL2) the increase is more significant than when the network power consumption is higher (PON). Introducing reduced power states in the network (IV) further increases the energy efficiency. These improvements are more significant with larger differences between P^n_0 and P^n_{red}. When applying all optimizations the thin client scenario is more energy efficient than the desktop scenario for passive user fractions up to approximately 95%, for all network technologies.

Between ADSL2, VDSL2 and PON there is a trade off. When all users are active ADSL2 is clearly more advantageous. However, for $N^u_{off}/N_u > 97\%$ PON is more efficient due to the large gap between the P^n_0 and P^n_{red}. This is displayed in Fig. 5. This trade off will be important for implementations where large passive user fractions during long periods of time can be expected. In this case for $\frac{N^u_{off}}{N^u} > 97\%$.

The passive state analysis clearly shows that the choice of a low power network technology with the possibility of reduced power states is required in order to assure power efficiency even with a large number of passive users.

5. Conclusions

ICT represents a relevant fraction of the worldwide energy production. The growth rate of this fraction is difficult to sustain. We created an analytical model in order to determine if the thin client paradigm is more power efficient than the desktop PC. Using experimental data different specific cases were reviewed.

These cases displayed that power savings up to 300% are possible. However, this potential is impaired by a reduced efficiency when a fraction of the users is passive. This can be mitigated by selectively switching off servers when reduced activity occurs. Secondly, introducing reduced power states in the network make the thin client paradigm more power efficient for idle user ratio's up to 95%.

Further optimizations will be achieved by optimizing the power consumption of the servers in the data centers. This can be achieved by building more energy efficient machines and by improving the data centers Power Usage Effectiveness (PUE).

ACKNOWLEDGEMENTS

The work described in this paper was carried out with the support of the BONE project ("Building the Future Optical Network in Europe"), a Network of Excellence funded by the European Community's Seventh Framework. Part of the research leading to these results was done for the MobiThin Project and has received funding from the European Community's Seventh Framework (FP7/2007-2013) under grant agreement nr 216946.

References

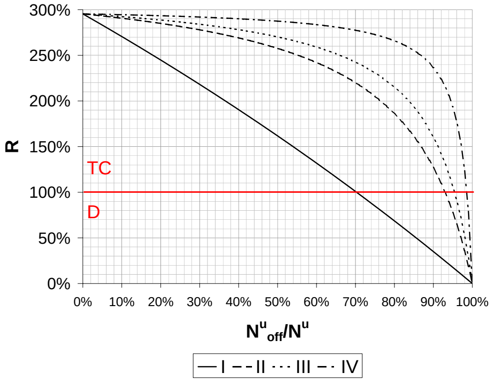

(a) ADSL2

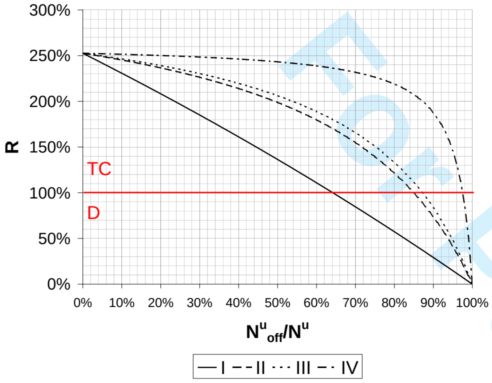

(b) VDSL2

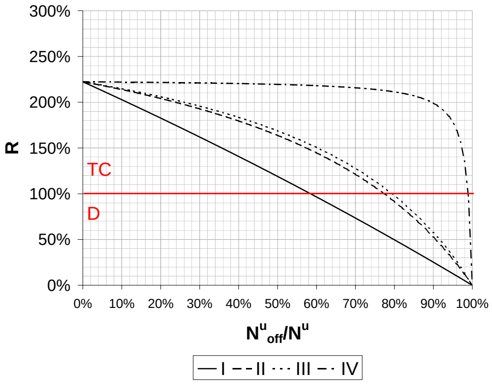

(c) PON

Figure 4. Power Saving Ratio of Thin Client towards Desktop PC in function of the fraction of passive users

1. J. Koomey, “Estimating total power consumption by servers in the U.S. and the world,” February 2007.
2. C. Cremers et al., “The impact of information and communication technologies on electricity consumption in germany,” in *ECEEE Summer Study*, ISBN 91-631-4001-2, St-Raphaël, France, June 2003, pp. 455–466.

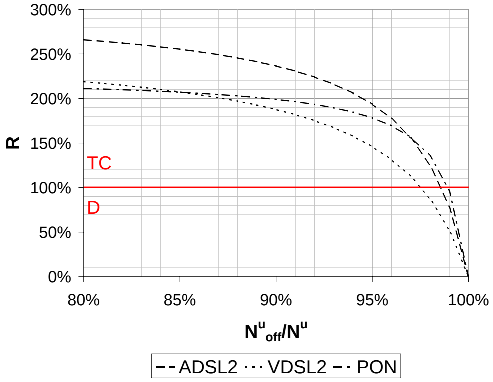

Figure 5. Power Saving Ratio of Thin Client towards Desktop PC using Reduced Network Power States

3. M. Pickavet, R. Van Caenegem, S. Demeyer, P. Audenaert, D. Colle, P. Demeester, R. Leppla, M. Jaeger, A. Gladisch, and H.-M. Foisel, “Energy footprint of ICT,” in *Broadband Europe 2007*, December 2007.
4. The green grid, http://www.thegreengrid.org.
5. IEEE 802.3 study group on energy-efficient ethernet, http://www.ethernetalliance.org.
6. C. Knermann, M. Hiebel, H. Pflaum, M. Rettweiler, and A. Schröder, “Environmental comparison of the relevance of PC and thin client desktop equipment for the climate,” *Fraunhofer UMSICHT*, 2008.
7. Standard Performance Evaluation Corporation (SPEC), http://www.spec.org.
8. “Green grid metrics: Describing datacenter power efficiency,” *The Green Grid*, 2007.
9. L. Deboosere, J. D. Wachter, P. Simoens, F. D. Turck, B. Dhoedt, and P. Demeester, “Thin client computing solutions in low- and high-motion scenarios,” in *Proc. ICNS 2007*, June 2007.
10. A.-C. Orgerie, L. Lefèvre, and J.-P. Gelas, “How an experimental grid is used: The Grid5000 case and its impact on energy usage,” *Proc. 8th IEEE Int. Symp. on Cluster Computing and the Grid (CCGrid2008)*, May 2008.
11. European Commission, “Code of conduct on energy consumption of broadband equipment version 2,” July 2007.

DOI: 10.1002/ett

AUTHORS' BIOGRAPHIES

WILLEM VEREECKEN received his M. Sc. degree in Electrical Engineering from Ghent University, Belgium, in July 2005. He is a Ph.D. student affiliated with the Department of Information Technology of Ghent University. His research is funded by a Dehousse Ph.D. grant from Ghent University. His main research interest is the design of energy efficient network architectures.

LIEN DEBOOSERE received her M. Sc. degree in Computer Science Engineering from Ghent University, Belgium, in July 2005. She is a Ph.D. student affiliated with the Department of Information Technology of Ghent University. Her research is funded by a Ph.D. grant from the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT Vlaanderen). Her main research interest is the design of architectures for wide area mobile thin client computing.

PIETER SIMOENS received his M. Sc. degree in Electronic Engineering from Ghent University, Belgium, in July 2005. Since October 2005, he is affiliated as research assistant with the Department of Information Technology of Ghent University. He is funded through the Fund for Scientific-Research-Flanders, Belgium (F.W.O.-V.). His main research interest is the development of intelligent and adaptive network architectures and protocols designed for thin client computing and autonomous environments. He cooperated in the IST FP6 project MUSE and is now work package leader for the protocol design in the IST FP7 project MobiThin.

BRECHT VERMEULEN received his Electronic Engineering degree in 1999 from the Ghent University, Belgium. In June 2004, he received his PhD degree for the work entitled 'Management architecture to support quality of service in the internet based on IntServ and DiffServ domains.' at the department of Information Technology of Ghent University. Since June 2004, he is leading a research team within the IBCN group of Prof. Demeester which investigates network and server performance and quality of experience in the fields of video, audio/voice and multiple play. Since the start of IBBT (Interdisciplinary institute for BroadBand Technology) in 2004, he leads also the IBBT Technical Test centre iLab.t in Gent, Belgium.

DIDIER COLLE received a M. Sc. degree in electrotechnical engineering from the Ghent University in 1997. Since then, he has been working at the same university as researcher in the department of Information Technology (INTEC). He is part of the research group INTEC Broadband Communication Networks (IBCN) headed by prof: Piet Demeester. His research lead to a Ph.D degree in February 2002. He was granted a postdoctoral scholarship from the Institute for the Promotion of Innovation through Science and Technology in Flanders (IWT-Vlaanderen) in the period 2003-2004. His research deals with design and planning of communication networks. This work is focussing on optical transport networks, to support the next-generation Internet. Up till now, he has actively been involved in several IST projects (LION, OPTIMIST, DAVID, STOLAS, NOBEL and LASAGNE), in the COST-action 266 and 291 and in the ITEA/IWT TBONES project. His work has been published in more than 100 scientific publications in international conferences and journals.

CHRIS DEVELDER received the M.Sc. degree in computer science engineering and a Ph.D. in electrical engineering from Ghent University (Ghent, Belgium), in July 1999 and December 2003 respectively. From October 1999 on, he has been working in the Department of Information Technology (INTEC), at the same university, as a Researcher for the Research Foundation-Flanders (FWO), in the field of network design and planning, mainly focusing on optical packet switched networks. In January 2004, he left University to join OPNET Technologies, working on transport network design and planning. In September 2005, he re-joined INTEC at Ghent University as a post-doctoral researcher, and as a post-doctoral fellow of the FWO since October 2006. Since October 2007 he holds a part-time professor position at Ghent University. He was and is involved in multiple national and European research projects (IST Lion, IST David, IST Stolas, IST Phosphorus, IST E-Photon One, BONE). His current research focuses on dimensioning, modeling and optimizing optical Grid networks and their control and management, as well as multimedia and home network software and technologies. He is an author or co-author of over 50 international publications.

MARIO PICKAVET received an M.Sc. and Ph.D. degree in electrical engineering, specialized in telecommunications, from Ghent University in 1996 and 1999, respectively. His current research interests are related to broadband communication networks (WDM, IP, (G-)MPLS, Ethernet, OPS, OBS) and include design, long-term planning, techno-economical analysis and energy efficiency of core and access networks. Special attention goes to Operations Research techniques that can be applied for routing and network design. In this context, he is currently involved in several European and national projects, such as the Network of Excellence "Building the Future Optical Network in Europe" (BONE). He has published about 200 international publications, both in journals (IEEE JSAC, IEEE Comm. Mag., Journal of Lightwave Technology, Eur. Trans. on Telecommunications, Photonic Network Communication, ...) and in proceedings of conferences. He is one of the authors of the book 'Network Recovery: Protection and Restoration of Optical, SONET-SDH, IP, and MPLS'.

BART DHOEDT received a degree in Engineering from Ghent University in 1990. In September 1990, he joined the Department of Information Technology, Ghent University. His research, addressing the use of micro-optics to realize parallel free space optical interconnects, resulted in a Ph.D. degree in 1995. After a 2 year post-doc in opto-electronics, he became professor at the Department of Information Technology. He is responsible for courses on algorithms, programming and software development. His research interests are software engineering and mobile & wireless communications. His current research addresses software technologies for communication networks, peer-to-peer networks, mobile networks and active networks.

PIET DEMEESTER received the Masters degree in Electro-technical engineering and the Ph.D degree from Ghent University, in 1984 and 1988, respectively. In 1992 he started a new research activity on broadband communication networks resulting in the IBCN-group (INTEC Broadband communications network research group). Since 1993 he became professor at Ghent University where he is responsible for the research and education on communication networks. The research activities cover various communication networks, including network planning, network and service management, telecom software, internetworking, network protocols for

QoS support, etc. He is member of the editorial board of several international journals and has been member of several technical program committees (ECOC, OFC, DRCN, ICCCN, IZS).

DOI: 10.1002/ett