Energy Efficiency in Thin Client Solutions

Willem Vereecken, Lien Deboosere, Pieter Simoens, Brecht Vermeulen, Didier Colle, Chris Develder, Mario Pickavet, Bart Dhoedt and Piet Demeester

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

Summary. In current society it is becoming more and more important to take energy efficiency considerations into account when designing solutions. In ICT, virtualisation is being regarded as a way to increase energy efficiency. This paper analyses the energy saving opportunities of the thin client paradigm.

1 Introduction

The current image of ICT is rather environmentally friendly. The worldwide communication via datacom and telecom networks has transformed society and created opportunities to reduce global energy consumption and CO_2 emissions in general. However, the ubiquitousness of ICT in daily life has caused its share in the global energy consumption to increase drastically. It is to be expected that this share will grow even more in the coming years. 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% [1].

Currently the power saving solutions for ICT were based on the principle of downscaling the performance of devices and even shutting them down when possible. A good example is mobile computing where devices need to be power efficient in order to maximize battery lifetime. On the other hand, power can be saved by assuring that a certain task is performed on the location where it will consume the least ammount of energy.

The power saving potential of this solution can be analysed with the thin client paradigm[2]. This approach is similar to 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.

Thin client solutions are currently implemented in order to reduce equipment cost and increase manageability. In this paper, however, we will analyse the implications of the thin client paradigm on power consumption at the customer premise, in the network and in the data centre. Based on this analysis we will try to determine the key aspects to consider when designing a power efficiënt thin client solution.

2 Mathemathical model

In order to determine the energy efficiency we will compare the power consumption of a standalone desktop with the power consumption of a thin client solution. For the thin client solution we consider the power consumption at the user premises, in the access network and in the data centre. These cases are schematically depicted in Fig. 1.

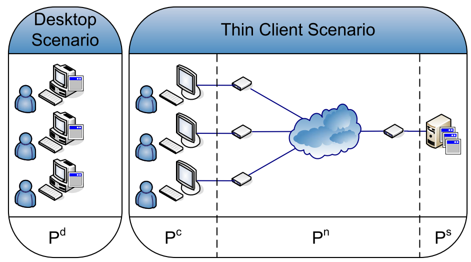

Fig. 1. Desktop and Thin Client Scenario

Note that we are fully allocating the power consumption in the network to the thin client solution. This is first to set clear boundaries for the analysis. Secondly, the thin client paradigm will be responsible for the majority of the traffic between the client terminal and the server ($0Mb/s - 5Mb/s$[3]).

2.1 Desktop, Client Terminal and Server

We will consider a linear model for the power consumption of a desktop computer (d), a thin client terminal (c) and a server (s). This power consumption will depend on the CPU load for running the application, denoted as λ^*_{CPU}, a number between 0 and 100%. The influence of the network traffic on the power consumption is negligable for the considered bandwidth. Thus, the model for a computer is (* = d,c,s):

$$P^* = P^*_0 + \alpha^*_{CPU}\lambda^*_{CPU} \tag{1}$$

For the client terminal, the power consumption appears to be constant even with varying CPU load λ^c_{CPU}.

On the server, we need to determine the dependency between λ^s_{CPU} and λ^d_{CPU}. Every calculation that needs to be performed on the desktop computer, needs to be performed on the server. Moreover, on the server there is also an overhead of the thin client protocol.

In order to be able to compare the CPU's on both the desktop and the server we denote the processing capacity of a server (according to a relevant performance oriented benchmark such as SPEC CINT2006 [4]) as C^s and the analogous parameter for the desktop case C^d. Since SPEC CINT2006 is a single threaded benchmark, we define the processing capacity as:

$$C^* = \#cores \times CINT2006 \tag{2}$$

We denote ϵ as the extra load per user caused by the thin client protocol. When we assume a share ratio of N users per server, the the CPU load on the server is:

$$\lambda^s_{CPU} = N\left[\lambda^d_{CPU}\frac{C^d}{C^s} + \epsilon\right] \tag{3}$$

This leads to the following relations:

$$P^d = P^d_0 + \alpha^d_{CPU}\lambda^d_{CPU} \tag{4}$$
$$P^c = P^c_0 \tag{5}$$
$$P^s = P^s_0 + N\alpha^s_{CPU}\left[\lambda^d_{CPU}\frac{C^d}{C^s} + \epsilon\right] \tag{6}$$

In [5], [6], [7] and [8] we can find data on the power consumption of desktops, laptops, servers and thin client devices respectively. In [4] we find reports with the CINT2006 benchmark.

2.2 The network

In order to limit the network latency we assume the data center to be located in the access network. [9] mentions target values for the power consumption of the network equipment.

2.3 Cooling

The servers are located in a data center. In a data center we also need to account for power consumption of HVAC, UPS, etc. This factor is denoted by the Power Usage Effectiveness (PUE) [10], the total power consumption of the data center, divided by the power consumption of the ICT equipment. Since our model should cover multiple cases we will consider the PUE accounted for in the relevant parameters.

2.4 Total Power Consumption of the Thin Client Solution

In summary, the power consumption for one user in the thin client paradigm is:

$$P^{tc} = P^c_0 + P^n + \frac{P^s_0}{N} + \alpha^s_{CPU}\left[\lambda^d_{CPU}\frac{C^d}{C^s} + \epsilon\right] \tag{7}$$

3 Equipment Selection

3.1 Desktop PC

In Fig. 2 we have the values for P_0^d and α_{CPU}^d for the category A, B and C computers from [5] and [6]. We see there is a strong correlation between P_0^d and α_{CPU}^d. For a desktop PC we can roughly say:

$$\frac{\alpha_{CPU}^d}{P_0^d + \alpha_C^d PU} = 72.7\% \tag{8}$$

For a laptop PC we get:

$$\frac{\alpha_{CPU}^d}{P_0^d + \alpha_C^d PU} = 67.9\% \tag{9}$$

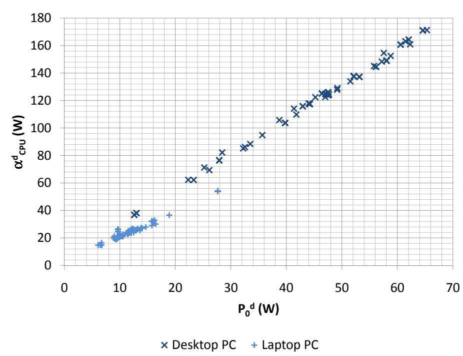

Fig. 2. Power consumption of Desktop and Laptop PC's

The power consumption of a laptop is also significantly lower than the power consumption of a desktop. However, this is not a fair comparison. The laptop PC performs exactly the same functionality as the desktop PC while only consuming a fraction of the power. This is because the laptop PC is optimized for maximal battery lifetime. This is not the case for all other devices used in the thin client solution. Therefore we want to compare technologies which are on the same level of power efficiency while clearly indicating that power optimizations of the involved equipment and an improved PUE in the data center will be required for the thin client paradigm to become a power efficient technology.

Based on these results we have selected the desktop *Dell OptiPlex360 (Intel Core 2 Duo E7400)* as a reference desktop computer. Its power consumption and processing capacity are summerized in table 1(a).

The average load λ_{CPU}^d will be approximately 10% on the desktop PC which is largely sufficient for standard office applications such as text editors and spreadsheets.

3.2 Thin Client Terminal

First we consider the power consumption of the client terminal. [8] mentions power consumption data for client terminals. This data is presented in Fig. 3.

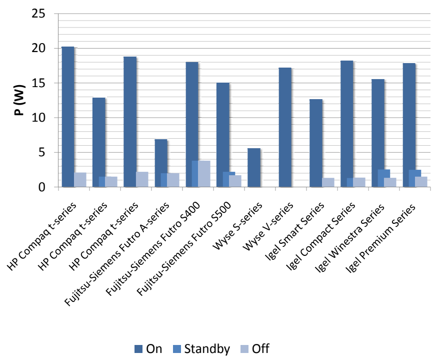

Fig. 3. Power Consumption of Client Terminals

For most devices the power consumption is comparable to that of a laptop PC. This is due to the amount of processing capacity and other functionality on the device. In some cases one cannot speak of a 'thin' client anymore and the term 'lean' client is used. In this study we want the capacity of the client device to be limited to only input-output signals. Therefore we use a *Wyse S10* device.

3.3 Server

For the servers there is less correlation between P_0^s and α_{CPU}^s. Moreover, we will try to have a maximal number of users on each server. This means $\lambda_{CPU}^s = 1$.

When we want to have an energy efficient solution we want the processing capacity per consumed power to be as high as possible. In Fig. 4 both values are given. [7] provides us with power consumption data. [4] provides us with the CINT2006 benchmark.

Generally speaking the power consumption scales with growing capacity. This is logical since C^s scales with the number of cores. There are however some servers which demonstrate a high capacity compared to the power consumption. Therefore we select a *ASUSTeK Computer ASUS RS160-E5 (2 × Intel Xeon L5420 Processor, 2,50 GHz)*. Its power consumption and processing capacity are summerized in table 1(a). The server overhead ϵ of the thin client protocol is considered to be small ($\epsilon \approx 0$). We also assume the server to be located in a data center with a typical PUE of 2.

3.4 Network

Finally, for the network power consumption we will base the used values on the target values mentioned in [9]. We consider three network technologies: ADSL2,

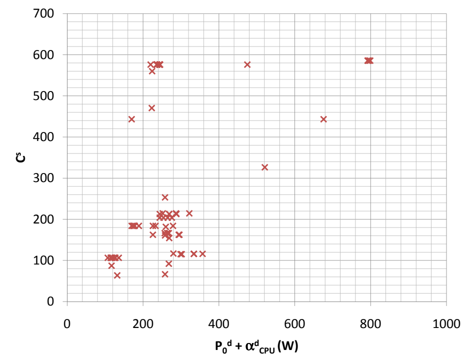

Fig. 4. Processing Capacity of Servers in function of their Maximal Power Consumption

Table 1. Equipment Parameters

(a) Power Consumption of Computers

Desktop PC	P_0^d	$39.764W$
	α_{CPU}^d	$103.736W$
	C^d	2×23
Client Terminal	P_0^c	$5.6W$
Server	P_0^s	$65W$
	α_{CPU}^s	$155W$
	C^s	16×36
	PUE	2

(b) Power Consumption per User of Network Equipment

	User Prem. Eq.	Access Netw. Eq.	Total
	Active state		
ADSL2	3.8 W	1.2 W	5.0 W
VDSL2	6.0 W	1.8 W	7.8 W
PON	7.7 W	11.0 W/32	8.04 W
	Reduced Power State		
ADSL2	2.6 W	0.4 W	3.0 W
VDSL2	3.5 W	0.6 W	4.1 W
PON	4.0 W	0.0 W	4.0 W

VDSL2 and PON. The network power consumption values are summarized in table 1(b).

4 Active State Analysis

Fig. 5 displays a breakdown in the power consumption for a Desktop PC and a Thin Client Setup. We have assumed the maximal share ratio of 125 on the servers.

Compared to the Desktop PC the power consumption of Thin Client Setup is significantly lower. We also notice that the power consumption of the thin client solution does not contain a dominant factor. This means that power optimizations at user premises, in the network and in the data center are equally important. For example, one would expect the PON solution to be most power efficient, but due to the high power consumption of the local gateway, this advantage is lost.

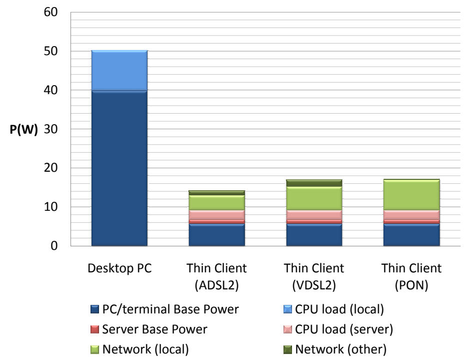

Fig. 5. Power consumption of Desktop PC and Thin Client in active state for $\lambda^d_{CPU} =$ 10% and N=125

We evaluate the power saving ratio $R = \frac{P_d}{P_{tc}}$ which expresses the relative power saving between both scenarios. The criterium for power efficiency is:

$$R > 100\% \tag{10}$$

In Fig. 6 R is displayed in function of the server share ratio N. We see the thin client solution is already very power efficient with low share ratio's. On the other hand, we see power saving ratio's up to 350% are achievable. Note however that this can impaired by the power consumption of the network.

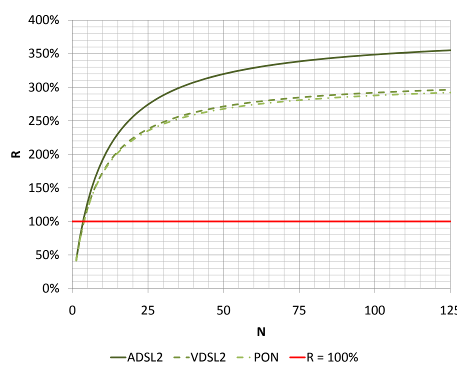

Fig. 6. Power Saving Ratio in function of the Server Share Ratio

5 Passive State Analysis

Until now we have assumed that all users are active. This is however not always the case. In this section we investigate the influence of passive users. We denote the number of active users as N^u_{act} and the number of passive users as N^u_{off}. Obviously, we always have:

$$N^u = N^u_{act} + N^u_{off} \tag{11}$$

For the desktop, client terminal and server we will assume that the device is either active or switched off. When a device is switched off it means it can be physically cut off from its power supply. In reality, this is not always the case and often there is a (low) standby power consumption. However, since we are aiming for a power efficient solution we will assume we cut off the power when a device is switched off.

In the network we do not cut off the devices since we want to keep a minimal connectivity between the user premise and the data center in order to be able to send wake-up signals to the devices. The reduced power state power consumption is given in table 1(b).

In the desktop scenario the power consumption will scale with the number of active users since all equipment is switched off for a passive user:

$$P^d_{PS} = \frac{N^u_{act}}{N^u} P^d_{act} = \left(1 - \frac{N^u_{off}}{N^u}\right) P^d \tag{12}$$

With P^d given by (4)

In the thin client scenario this is less straightforward.

In Fig. 7 the influence of the number of passive users, which have their client terminal switched off, on the power saving ratio is displayed. We have evaluated the relevant scenario's in the ADSL2 case. For λ^d_{CPU} we assumed a value of 10%. We assumed the maximal share ratio of $N = 125$.

When we do not mitigate for the passive users and all the servers remain active the efficiency degrades approximately linearly in function of the fraction of passive users $\frac{N^u_{off}}{N^u}$. This can be explained because P^d_{PS} degrades linearly with a rising passive user fraction while $P^t c_{PC}$ almost remains constant. The only factor reducing the power consumption is the switched off client terminal.

We can however measure the number of active users and only switch on the required number of servers so that the active servers are used at their full capacity and the passive servers are switched off. This already leads to a significant optimization.

Finally we also use a reduced power state to connect the passive users to the data center which leads to an even more optimized power consumption.

It is clear that the optimization solutions allow for an increasing number of passive users to keep the thin client solution more efficient than the desktop solution.

6 Optimization parameters for power consumption

In the previous sections we have elaborated on the power saving potential of the thin client paradigm by modelling a power efficient solution. In this section

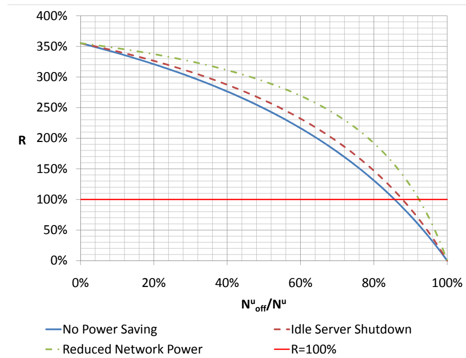

Fig. 7. Power Saving Ratio of ADSL2 case in function of the fraction of passive users

we will investigate the influence of the different equipment parameters on this potential.

Therefore we define two criteria by which we evaluate a thin client solution. The first is the maximal power saving ratio R_{max}. This is the power saving ratio when all users are active. The higher R_{max}, the better the solution.

The second criterium is the break-even passive user fraction $\left[\frac{N^u_{off}}{N^u}\right]_{R=1}$. This is the passive user fraction for which the power saving ratio is 100% or the thin client solution is as efficient as the desktop solution.

In Fig. 8 we have displayed both criteria in function of the maximal share ratio of the servers. We based our data on [7]. Note that the maximal share ratio is proportional to the capacity of the server (cfr. (3)).

We see that a higher share ratio often implies that a high capacity can be reached. Power saving ratio's of approximately 350 % are possible. Higher share ratio's also have a positive effect on the break-even passive user fraction and fractions up to 90% are possible when we do no power saving. This effect obviously diminishes when mitigating through idle server shutdown.

We also look at the influence of the power consumption of the client terminal and the network. This is displayed in Fig. 9. It is clear that both have a large impact on the power efficiency and the selection of an energy efficient client terminal and network technology will be important.

7 Conclusions

In this paper we created an analytical model in order to investigate the power efficiency of the thin client paradigm.

Comparing the paradigm with a laptop PC has shown that power optimizations of the individual equipment and the datacentre PUE will be required. However, when comparing with technology with a similar level of energy optimization (Desktop PC) the thin client paradigm shows a clear potential.

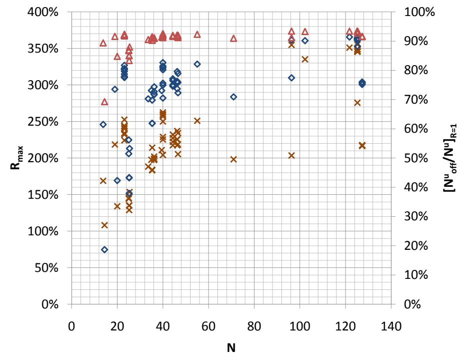

Fig. 8. Influence of the maximal server share ratio on R_{max} ($\times$) and $\left[\frac{N^u_{off}}{N^u}\right]_{R=1}$ with No Power Saving ($\diamond$) and Idle Server Shutdown ($\triangle$)

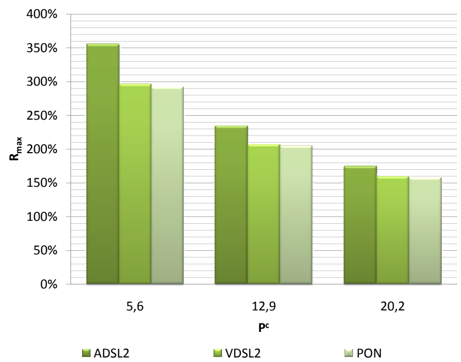

Fig. 9. Influence of the client terminal power consumption and network technology on R_{max}

The cases displayed that power savings up to 350% 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 97%.

It is important to select servers with a high capacity so a large number of sessions can run on it. Secondly, it is important to select a low power client device and an energy efficient network to interconnect the device with the servers.

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

1. M. Pickavet, W. Vereecken, S. Demeyer, P. Audenaert, B. Vermeulen, C. Develder, D. Colle, B. Dhoedt, and P. Demeester, "Worldwide energy needs for ICT: the rise of power-aware networking," in *IEEE ANTS 2008*, December 2008.
2. C. Knermann, M. Hiebel, H. Pflaum, M. Rettweiler, and A. Schroder, "Environmental comparison of the relevance of pc and thin client desktop equipment for the climate," *Fraunhofer UMSICHT*, 2008.
3. L. Deboosere, J. De Wachter, F. De Turck, B. Dhoedt, and P. Demeester, "Thin Client Computing Solutions in Low- and High-Motion Scenarios," in *International Conference on Networking and Services (ICNS)*, 2007, p. 38.
4. Standard Performance Evaluation Corporation, "Spec cpu2006," http://www.spec.org/cpu2006/, July 2009.
5. Energy Star, "Qualified desktops and integrated computers," http://www.energystar.gov, July 2009.
6. Energy Star, "Qualified notebooks and tablet computers," http://www.energystar.gov, July 2009.
7. Standard Performance Evaluation Corporation, "Specpower_ssj2008 results," http://www.spec.org/power_ssj2008/results/, July 2009.
8. The Danish Electricity Saving Trust, "See the trust's data covering power used by different thin clients," http://www.savingtrust.dk/, March 2009.
9. European Commission, "Code of conduct on energy consumption of broadband equipment version 3," November 2008.
10. T. G. Grid, "Green grid metrics: Describing datacenter power efficiency," http://www.thegreengrid.org, February 2007.