A

Green IT - an analysis of available data and guidelines for reducing energy consumption

LUCA ARDITO, Politecnico di Torino
MAURIZIO MORISIO, Politecnico di Torino

Green IT has become one of the most discussed topics in the last new years. However, the related literature is dispersed and it is difficult to give precise definitions. Nowadays saving energy is an interesting and interdisciplinary key challenge. Hardware manufacturers and designers have usually handled the problem, in the field of IT, but recently software energy efficiency gathered the interest of industry and academic research. However, writing energy efficient software requires proper metrics to evaluate it. In the literature it is possible to find metrics related to several aspects of software evaluation such as complexity, performance, maintainability, reliability, and so on but there is still a gap for energy related metrics. In this paper we:

— Introduce a taxonomy of concepts related to energy and IT;
— Present the most recent data on energy consumption trends organized according to the taxonomy;
— Present some guidelines to write energy efficient software organized according to the taxonomy.

Our contribution is twofold:

— Provide available information in a better organized way;
— Underline what is missing and what can be done to make the context clearer as well.

Additional Key Words and Phrases: Green IT, Power management, Energy Efficiency, ICT Energy Reduction, Software Energy Consumption

1. INTRODUCTION

The rapid growth and significant development of IT systems has started to cause an increase of worldwide energy consumption [Webb and Al. 2008]. This issue moved technology producers, information systems managers, and researchers to deal with energy consumption reduction in terms of:

— Global CO2 footprint;
— Consumption of data centers;
— Reduced battery life of portable devices;
— Economic impact of a new business model, which aims at greening everything.

As well described in [Krikke 2008], IT producers are forced to manage the product lifecycle by legislations, but also users are becoming concerned about the environmental implications from the use of IT. This area of research is called GREEN IT. GREEN IT "refers to environmentally sustainable computing or IT [1] " and is defined as "The study and practice of designing, manufacturing, using, and disposing of computers, servers, and associated subsystems such as monitors, printers, storage devices, and networking and communications systems efficiently and effectively with minimal or no impact on the environment" [Murugesan 2008]. Murugesan in [Murugesan 2008] also expresses the benefits introduced by Green IT: "Green IT benefits the environment by improving energy efficiency, lowering greenhouse gas emissions, using less harmful material, and encouraging reuse and recycling". Taking into account the research areas identified in [Capra and Merlo 2009] and according to [Murugesan 2008], the main topics related to Green IT, are:

— Design for environmental sustainability: "balancing energy and resource savings by ICT and energy and resource consumption of ICT" [Naumann et al. 2011] and "making business operations, buildings, and other systems energy efficient" [Murugesan 2010];

[1] http://en.wikipedia.org/wiki/Green_computing Last Visited: 23 January 2012

— Energy Efficient Computing: the efficient use of resources in terms of energy-aware algorithms;
— Power Management: a set of HW/SW techniques that optimize the management of power resources in computer systems, portable devices, and data centers.
— Data center design: eco-friendly devices that improve energy efficiency and energy conservation of data centers (i.e., energy-efficient mechanical and electrical systems, green power, use of natural light, etc.);
— Virtualization: "the faithful reproduction of an entire architecture in software, which provides the illusion of a real machine to all software running above it" [Kiyanclar 2005];
— Disposal and recycling management: managing e-waste, and limiting planned obsolescence upgrading devices instead of replacing them;
— Regulatory Compliance: Regulatory requirements and legislative actions tend to force acceptance of a technology or practice in situations where this would not occur. The existence of certain rules on sustainability in IT standards, can lead to the adoption of some green IT initiatives. [Molla 2009].
— Green metrics, tools and methodology assessments: software tools for collecting or simulating, analysing, modelling, reporting energy consumptions, environmental risk management, environmental impact, and greenhouse gas emissions; platforms for eco-management, emission trading, or ethical investing [Murugesan 2008];

Green IT involves many areas and stakeholders, starting from governments, through new business models and R&D, to different technical fields. IT can also be used to monitor energy consumption such as: heating systems in buildings, fuel efficiency in cars or smart grids implementation. So IT is involved in worldwide energy reduction, and in reducing its energy consumption as well. The main contribution of this work is to summarize evidence available in the literature about:

— Measuring techniques for energy consumptions in IT systems;
— IT energy consumption trend according to our taxonomy;
— Methodologies to reduce energy consumption of IT;

The goal is to offer a clearer picture of the state-of-the-art in this field, and to highlight areas where evidence is missing as well. This paper is organized as follows:

— Section 2 proposes our taxonomy.
— Section 3 deals with the energy metrics, which are used to characterize more quantitatively what "greenness" means at each layer of the taxonomy.
— Section 4 summarizes the current worldwide energy consumption and carbon footprint using our taxonomy.
— Section 5 describes some guidelines that can improve energy efficiency in IT systems.
— Section 6 gives our conclusions.

2. TAXONOMY

We define the following taxonomy to organize IT and energy consumption. We consider two orthogonal dimensions, the time axis (or IT energy lifecycle) and the space axis (IT elements and infrastructure).

2.1. IT Energy Life Cycle

Inspired by [Forge 2007] we propose the energy life cycle in Figure 1. The activities in the lifecycle are design of the IT product, manufacture, transport (includes packaging and distribution), use, disposal, and possibly recycle. All these activities consume materials and energy, with related emissions.

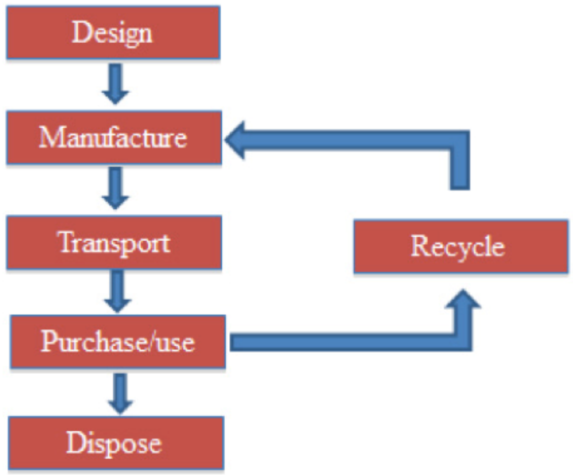

Fig. 1. The time axis: Energy Life Cycle

2.2. IT Elements and Infrastructure

On the spatial axis, we consider elements or nodes (e.g. PC, peripheral devices and mobile phones) and their connections (cables, wireless links etc.) to build infrastructures (e.g. PC networks, the Internet, data centers, the cloud). In a node, we define different layers (starting from hardware to application) as defined later (Figure 2)

- Network element: this element considers network equipment (e.g. Network Interface Card, router and gateways) and protocols, which means everything is related to connectivity.
- A node element consists of three layers:
  - Hardware layer: this layer considers CPUs, GPUs, and storage (memory, disks).
  - Operating System layer: this layer considers software programs implementing the traditional operating system services (file and memory management, task scheduling, I/O management). The key issue here is power management.
  - Application layer: this layer considers software to implement user level services. Key issues considered here are the energy efficiency of algorithms and software architectures.
- Infrastructure element: This element considers many nodes connected by a network to define a larger entity capable of offering higher-level computation services. Entities of this kind are data centers, web farms, and cloud computing.

3. ENERGY BENCHMARKS AND METRICS ASSIGNMENT

Energy (measured in Joule or Wh) and power (measured in J/s or W) are the metrics, which can be used to characterize consumption of IT and ICT systems. However, they are not specific to IT. In literature, other specific measures have been proposed. We can summarize them into three broad categories:

- Power, in terms of consumed Watts.
- Efficiency, as the ratio of useful energy and total energy used
- Productivity, defined, at high level, on a production process, as output/resource on a time interval (ex. cars produced per worker in a day). In the context of Energy and IT, the output is computational work while the resource is energy. Computational work needs to be defined at each level of the taxonomy. For instance: in a CPU, an example may be operations performed, in a network bits transmitted, in a web application hits managed.

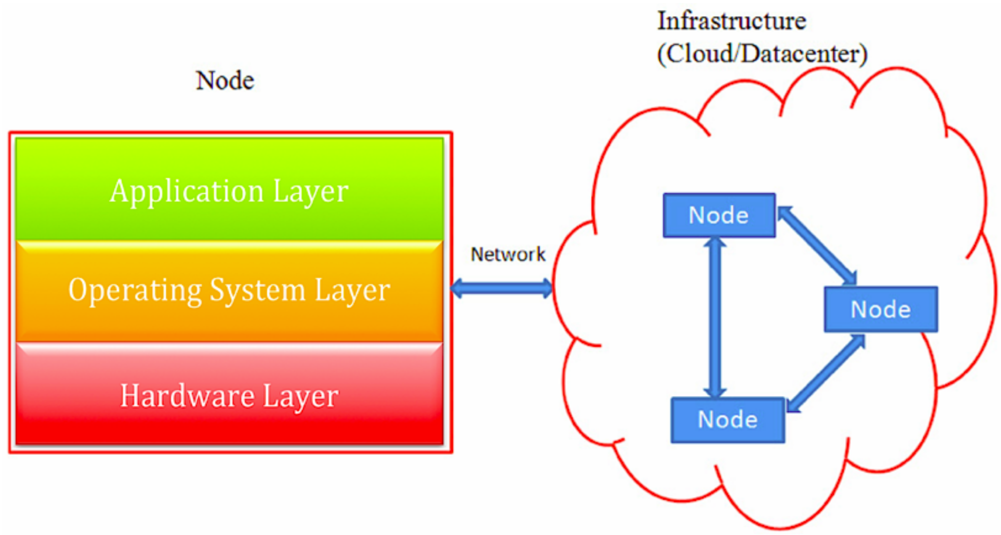

Fig. 2. The space axis: Nodes and Infrastructure

Table I reports measures proposed in literature [Rivoire et al. 2007b], [Bianzino et al. 2011] and [Gude and Lago 2010] placing them in our taxonomy and categorizing them as efficiency or productivity. Not all layers of the taxonomy are covered, meaning that other efficiency or productivity measures could be defined.

4. ENERGY CONSUMPTION AND CARBON FOOTPRINT

A lot of data has been published on IT related consumption and emissions. However these data are usually sparse. In this section we summarize them using the taxonomy as a unifying point of view. In 4.1 we review data at a global level, then we will focus on the space and time view (4.2, 4.3), and further down.

4.1. The global view

To understand the impact of IT we need first to analyse the figures of primary energy consumption, electricity consumption and finally IT energy consumption. Besides, we need to analyse in the same way green house gases emissions. Let's start from an analysis of energy consumption. The annual global primary energy consumption in 2007, based on the International Energy Outlook 2007 published by United States Department of Energy [Staub 2010], was around 140,000 TWh. The distribution of consumption [Staub 2010] is in Figure 3.

Industry, namely, agriculture, mining, manufacturing, and construction, consumed about 37% of total energy. Personal and commercial transportation consumed 20%; residential heating, lighting, and appliances consumed 11%; and commercial uses (lighting, heating and cooling of commercial buildings, and provision of water and sewer services) consumed 5%. The other 27% of the world's energy is lost in energy transmission and generation. Since IT usage depends on electrical energy, we analyse now this aspect. The production of electrical energy in 2007 was around 17,400 TWh [Staub 2010] or the 12% of global primary energy consumption. Figure 4 reports the electrical energy consumption distribution in 2007, based on the report from IEA (International Energy Agency) [4]. Industry used 42%; residential used 27%; commercial and public services used 23%; and the electricity transportation used

[4] `http://tinyurl.com/bubgrlz` Last Visited: 7 May 2012

Table I. Energy Metrics and Benchmarks

Layer	Category	Unit	Description	Example
Infrastructure	Productivity	Useful work/W	Green Grid Datacenter Performance Efficiency (DCPE) [Grid]	Aims at measuring “Useful work” delivered by a data center vs. the power used
Infrastructure	Efficiency	%	Useful Power (for storage, computation, communication) / Total power	Green Grid Datacenter Efficiency(DCE) [Grid]
Node	Productivity	MFLOPS/W or FLOP/J	Number of Floating point operation computed per watt	The Green 500 list[2] ranks high performance computers on MFLOPS/W, instead of the usual ranking on FLOPS only
Application	Power	W	Power used by an application on a node	Joulemeter [Kansal et al. 2010] is a tool able to estimate instant power consumption of PCs and applications
Application	Productivity	Operation / J	Output is intended as sorted records	Joule Sort [Rivoire et al. 2007a], counts sorted records per Joule, and is a benchmark for sorting algorithms. It is a variant of Sort benchmark that considers records sorted per second
OS	Power	W	Power used by an OS or an app on a node	Softwatt [Gurumurthi et al. 2002] is a tool to estimate power used by OS and applications
HW	Power	W	SimplePower [3]	Given an instruction (or an instruction set) and a program, it estimates energy consumption on the CPU).
Network	Efficiency	% 100(MI)/M	I = energy consumption at idle, M energy at maximum	Environmental Performance Index (EPI) [Mahadevan et al. 2009]
Network	Productivity	KB/J	Bits transferred per Joule over a channel	Energy efficiency rating (EER) [Ceuppens et al. 2008] does the same in Gbps/W

the rest 8%, In the same year, the total energy consumption of the IT sector was estimated [7] at between 370-830 TWh, which means around 0.3-0.6% of the global primary energy or around 2-5% of the global electricity.

Considering carbon emissions, in 2007 the total carbon emission has been 830 MtCO2e[Kirby et al. 2008]. The distribution of these emission by origin is reported in Figure 5, based on the book: “Kick The Habit, A UN guide to climate neutrality” [Kirby et al. 2008] 64.70% of carbon emissions is related to industrial energy usage, transformation and transportation, e.g. industry, electricity, heat, and fuel combustion. This report does not provide an estimate of the IT impact. This figure is provided by [Webb and Al. 2008] as 2% of global carbon emissions.

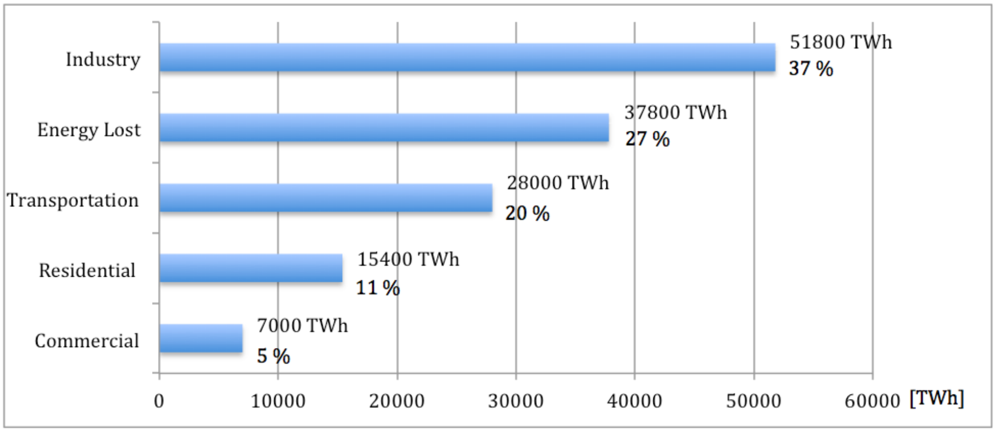

Fig. 3. Worldwide energy consumption distribution (2007)

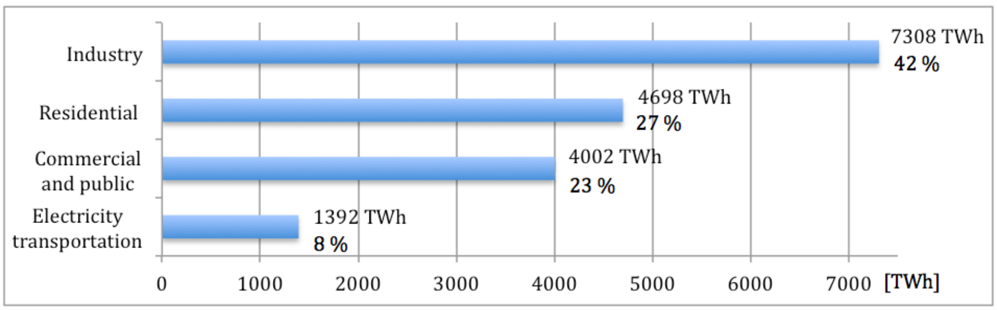

Fig. 4. Worldwide electricity consumption distribution (2007)

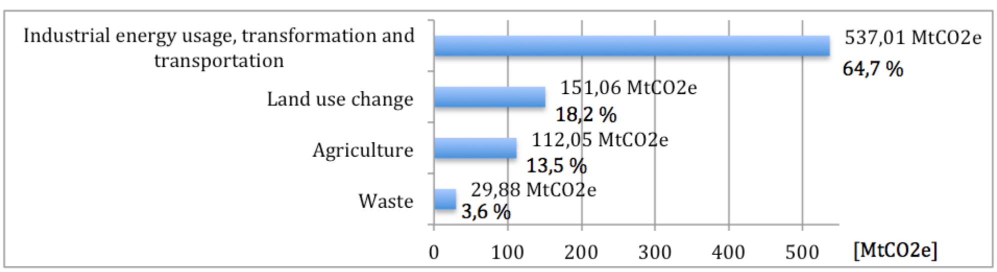

Fig. 5. CO2 emission distribution 2007

After this punctual analysis, let's consider future trends (2020). Total energy 180,000 TWh [5], electrical energy 27,400 TWh [PATEL, S 2011], and carbon emissions 1430 MtCO2e [MILES, T 2010]. In Figure 6 we report the carbon emission trend, based on [Webb and Al. 2008].

At this point we can conclude, even considering the unavoidable imprecision of the numbers presented, that, overall, IT is responsible for very limited percentages of consumptions and emissions as described in Figure 7 (less than 1% consumption of primary energy, around

[5]http://www.solcomhouse.com/worldenergy.htm Last Visited: 07 May 2012

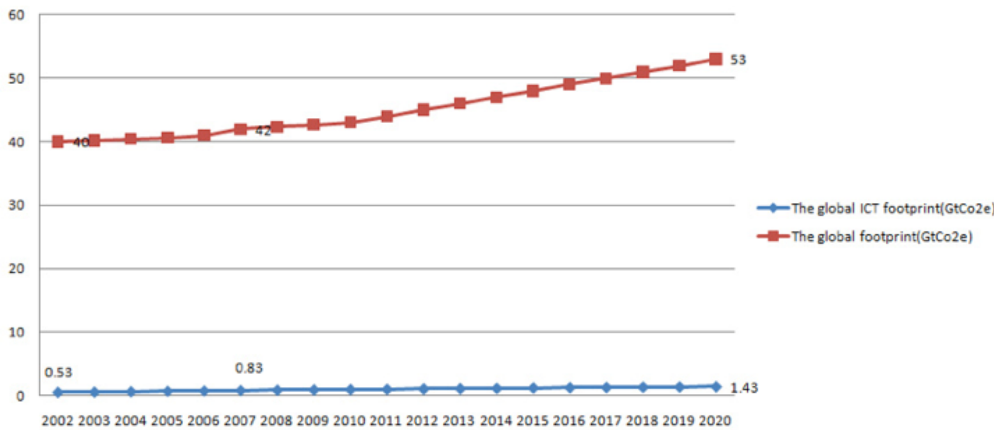

Fig. 6. Worldwide estimated CO2 footprint trend, global and for IT

3% of electrical energy, 2% of carbon emissions). However, an analysis of future trends shows that emissions and consumption tend to increase. At this level much work should be done to stabilize and possibly reduce these trends, in order to provide a sustainable IT future. The analysis performed before considers the whole of IT. In the following we analyse the problem considering the spatial and time dimensions, as introduced by the taxonomy.

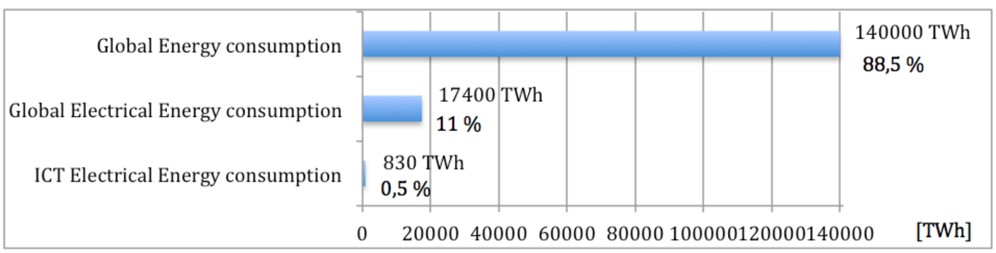

Fig. 7. Comparisons of world energy consumption, world electrical energy consumption, IT electrical energy consumption

4.2. Worldwide space dimension

Focusing on the IT sector, a key question is which IT components are responsible for these trends. We start from individual components (or nodes such as PCs, fixed and mobile, and data centers) and we conclude with networks.

4.2.1. Node view (PCs, Laptops, Mobile Devices) and Data Centers. According to [Webb and Al. 2008] we report some numbers about the number of different class of devices and their consumption from 2002 to 2020. The number of worldwide PCs is expected to grow from 592 million (2002) to more than 4 billion (2020) and, regarding energy consumptions, laptops will overtake desktop computers by 2020. In 2002, emissions of PCs and monitors were 200 MtCO2e, growing to 600 MtCO2e by 2020. The number of servers in 2002 was 18 million and there will be a sharp increase of this figure up to 122 million in 2020. In 2002 data center emissions were approximately equal to 76 MtCO2e and this value should more than triple by 2020 to 259 MtCO2e. In 2002 there were 1.1 billion mobile devices. This is expected to grow to 4.8 billion in 2020. Telecommunications emissions rise from 150 MtCO2e in 2002 to about 350 MtCO2e in 2020. Figure 8 represents the carbon emission in 2002 [Webb and Al.

2008], and Figure 9 reports an estimate, based on [Webb and Al. 2008] of carbon emissions by IT by 2020.

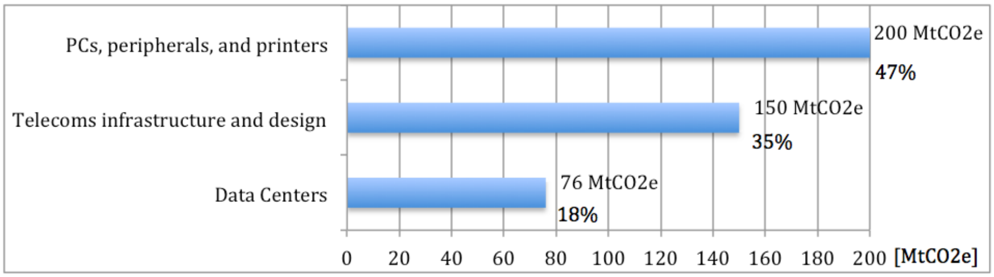

Fig. 8. The global footprint by subsector 2002

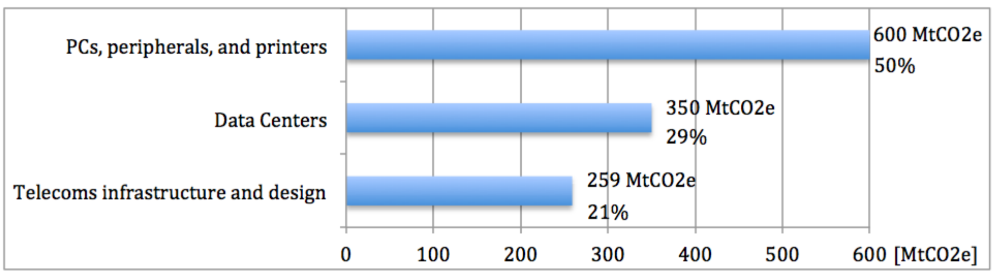

Fig. 9. The global footprint by subsector 2020

The majority (57 percent) will come from PCs, peripherals, and printers, while data centers account for 25% only. Figure 10 compares the growth of devices produced and the growth of carbon emissions [Webb and Al. 2008].

From the analysis above it is clear that PCs and mobile devices are the key actors in consumptions in the future. Another study [IDC 2009] stated that data centers in 2009 consumed about 330 TWh worldwide, and authors in [Somavat et al. 2010] estimated the electrical energy consumption of PCs, laptops, and mobile phones in 2009:

— PCs: 163,2 TWh
— Laptops: 46,2 TWh
— Mobile Phones: 44,6 TWh

These data are calculated in terms of electrical energy consumed and cannot be compared with data in Figure 8 because of different units, years, and data aggregations.

4.2.2. Network view. The number of Internet users has continuously increased, and the energy consumption of the Internet has grown accordingly. Network equipment such as hubs, switches and routers for the Internet consumed the energy of about 6.05 TWh/year in 2010 as shown in Figure 11 and it is expected to grow by 1 TWh or more per year.

In addition, the network equipment that connects to the Internet in home and office networks transmit packets via Ethernet links. The estimated energy consumption of Network Interface Cards (NICs) and other network devices, which use Ethernet links in the US, was approximately 5.3 TWh/yr in 2005 [Nordman and Christensen 2005]. Moreover, the default link rate of the Ethernet and the network edge devices is rapidly increasing from 10 Mbps to

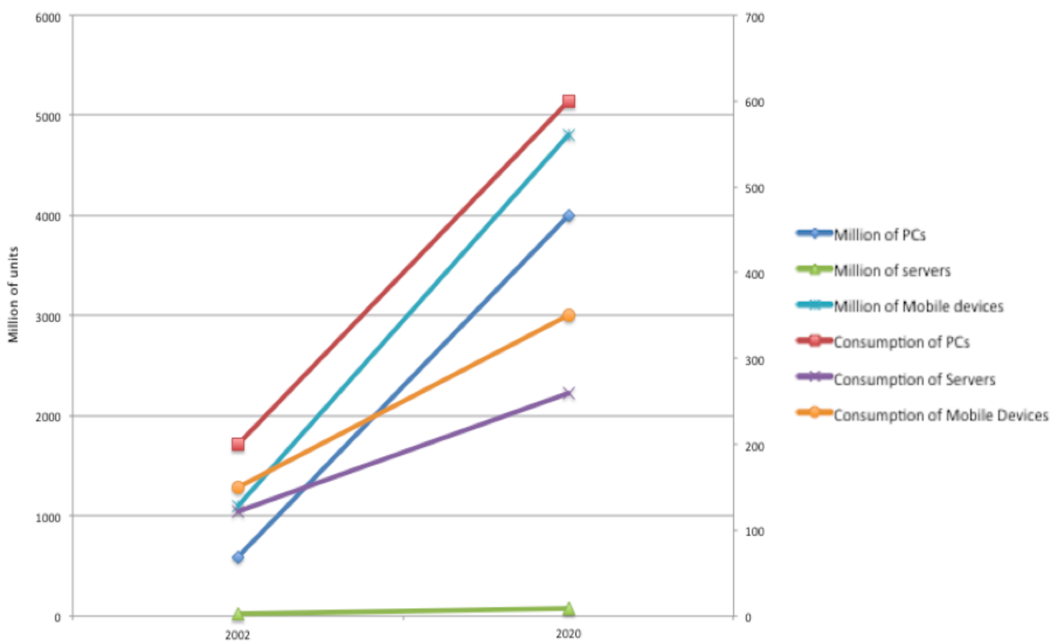

Fig. 10. PC, Servers, and Mobile Phones production vs. Carbon Emissions

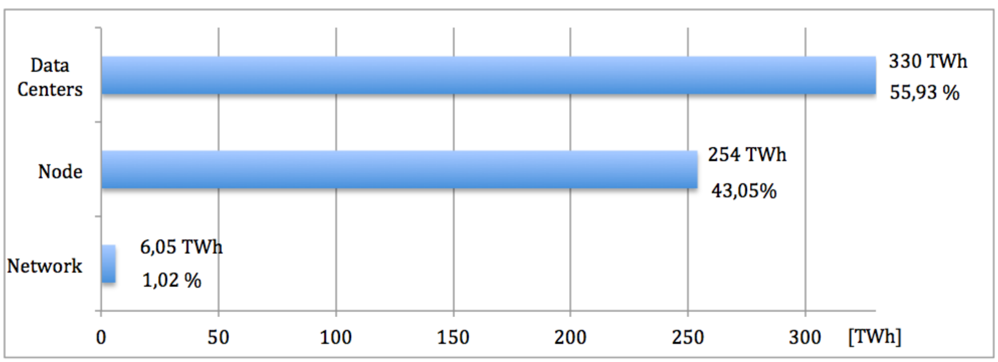

Fig. 11. Electrical Energy consumption of IT devices basing upon our taxonomy proposed in section 2.2

1 Gbps or more, and the number of network devices is also increasing [Bianzino et al. 2010]. To sum up, basing upon our taxonomy, most of the energy consumption is concentrated within data centers and nodes. Network devices have a less important role but not negligible because of the magnitudes involved.

4.3. Local space dimension

4.3.1. Within a Node. In Figure 12 we analyse consumption inside a node [Moshnyaga 2009]. The energy contribution due to software can be measured on hardware. This, from the hardware perspective, is seen as a different trend of the instant power consumed by the device. There are several (and canonical) ways to gather energy consumption data from a hardware device and we cite four examples. In 1998 authors of [Russell and Jacome 1998] stated that the current drawn by a processor could be measured using an oscilloscope with

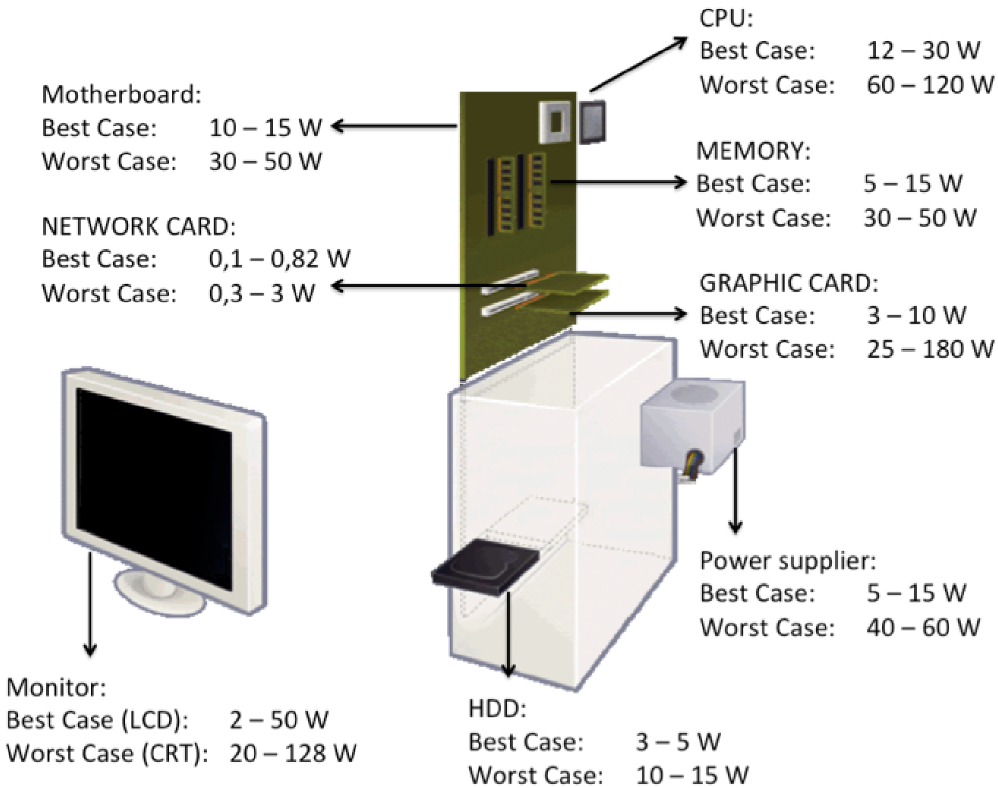

Fig. 12. Typical energy consumption of PC components based on [Moshnyaga 2009]

a shunt resistor, connected in series with the supply voltage pin of the microprocessor. It is also possible to perform a direct (current) measure by utilizing an ammeter with a high frequency signal to obtain a stable value [Tiwari and Lee 1995]. In [Kravets and Krishnan 1998] the authors connected a multimeter to a laptop power supplier. They also plugged the laptop supplier into a universal power supply (UPS) to filter out voltage fluctuations. Their multimeter sampled the current 11-12 times a second. Otherwise it is possible to use smart meters; some commercially available product are Plogg Meter [6], Kill-a-Watt / Tweet-a-Watt [7], and SmartLink [8]. They aim more at monitoring, controlling, and automating the attached device. It is very difficult to find reliable values about CPU power consumptions. Usually manufacturers publish the "Thermal design power" (TDP). This value does not match the maximum CPU power consumption but it refers to the maximum amount of power that the cooling system in a computer has to dissipate and the result is expressed in Watt. These values are not comparable between CPUs produced by different manufacturers. AMD introduced a new metric (called ACP) to measure the CPU power consumption [AMD 2010]. ACP is obtained by measurements taken on specially instrumented components in particular conditions (temperature, workloads, configurations). In [Jaiantilal et al. 2010] authors analysed the problem of estimating system power consumption. They stated that for each task, it is possible to measure the number of clock cycles executed per unit time

[6] `http://www.plogginternational.com` Last Visited: 23 Jan 2012
[7] `http://www.ladyada.net/make/tweetawatt/` Last Visited: 07 May 2012
[8] `http://www.smarthome.com` Last Visited: 18 Jan 2012

and generate a model that can predict the watt consumed: P(System) Power(Taski) + P(bias') Where bias' equals every other component and Power(Taski) = F * number of FP Cycles + I * number of Int Cycles + M *number of Memory Cycles. Authors in [Hylick et al. 2008] studied in deep server hard drives (3,5) energy consumption, both from a mechanical and electronic perspective. Based on their studies, drive platters spin constantly, they stop only in standby mode, read/write heads are only powered during the reading and writing phase. The arm actuator is only powered when there is the need to seek across locations on a platter. Printed circuit board electronics are instead always powered. Based on ATA/ATAPI-5 specification and the Advanced Configuration and Power Interface (ACPI), modern hard drives support four power management states: active, idle, standby, and sleep (it is not possible to recover from this state without a system reset). In standby mode, mechanical energy savings may range from 92 % to 99% and electronic components energy saving may span from 35% to 95%. In idle mode mechanical components can consume from 25% to 75% of the total energy consumption. During the read/write phase the energy consumption is dependent on the Logical Block Number (LBN). Data density increases at higher LBNs and more time is required to recover the data (due to constant angular velocity of the spindle), read bandwidths decreases and read energy consumption increases. But write bandwidths remains fairly constant because write requests are not influenced by the varying data density on the platters, hence bandwidth and energy consumption do not vary on the basis of LBN. In general, reading consumes more energy than writing for blocks larger than 2KB, and writing is more energy intensive than reading for blocks smaller than 2KB.

Table II. Energy Metrics and Benchmarks

Disk Model Unit	**Cap.** GB	**Read** MiB/s	**Write** MiB/s	**Read** W	**Write** W	**Idle** W	**Efficiency** MiB/j
SSD							
Intel X-25E	32	226	198	1.7	2.7	0.6	103
Intel X-25M	80	225	79	1.0	2.5	0.6	128
Samsung PB22-J	256	201	180	1.1	2.8	0.6	124
Super Talent FfM56GX25H	256	235	163	1.6	2.9	0.5	102
HDD							
Samsung HD502HI	500	106	108	6.6	6.6	3.7	16
Samsung HM500JI	500	87	87	2.3	2.3	/	28
W.D. WD7500KEV	750	82	82	2.0	2.0	/	41

The energy consumed during seeking is minimal, but restricting disk accesses between low and central LBNs could help to save energy in long usage periods. Solid-state drives have different power consumption profiles and efficiency as shown in Table II. Authors in [Roth et al. 2002] analysed energy consumption of office and telecommunications equipment in commercial buildings. They divided monitors and printers into categories. Table III and Table IV show the energy consumption of monitors divided into two categories: CRT and LCD. Over the years monitors show an increase of performance and a decrease of energy consumption.

Table III. CRT energy consumption estimate [Roth et al. 2002]

Type - CRT	**Active**	**Standby**	**Suspend**	**Off**
14-15"	61	53	19	3
17"	90	26	9.2	4.3
19"	104	31	13	4
21"	135	43	14	4.7

Table V shows the energy consumption of different categories of printers.

Table IV. LCD energy consumption estimate [Roth et al. 2002]

Type - LCD	Active	Standby	Suspend	Off
13"	2.5	0.7	0.2	0.1
14"	6.7	1.9	0.7	0.3
15"	11.7	3.4	1.2	0.6
17"	16.7	4.8	1.7	0.8
18"	25	7.2	2.5	1.2
20"	31.7	9.2	3.2	1.6
21"	36	10.4	3.6	1.8

Table V. Printer energy consumption estimate [Roth et al. 2002]

Device	Active	Standby	Suspend	Off
Impact Printer	36.5	16.8	N/A	1
Inkjet Printer	42.576	13.377	N/A	2.878
Laser Printer	231	28	16	1.9
Laser Printer Small Desktop	130	75	10	N/A
Laser Printer Desktop	215	100	35	N/A
Laser Printer Small Office	320	160	70	N/A
Laser Printer Large Office	550	275	125	N/A

4.3.2. Within Network, Data Centers and Infrastructure. Authors in [Gupta and Singh 2003] examined the energy consumption of networking devices in the Internet, basing upon data collected by the U.S. Department of Commerce. Data are available in Table VI.

Table VI. Energy consumption of networking devices in the Internet [Gupta and Singh 2003]

Device	Approximate Number Deployed	Total AEC TW-h
Hubs	93.5 Million	1.6 TW-h
LAN Switch	95,000	3.2 TW-h
WAN Switch	50,000	0.15 TW-h
Router	3,257	1.1 TW-h
Total		6.05 TW-h

Basing upon these findings, the impact given by a reduction in consumption in this field can be relevant.

4.4. Time dimension

4.4.1. PC. Let's now analyse how these consumptions are distributed over the lifecycle of IT devices (manufacturing and transport, usage, disposal). According to [Williams 2004] the energy to manufacture a PC accounts to 4250 MJ, the energy spent in usage (considering an average usage time of 3 years) is 1500 MJ, and the overall energetic cost (including transport and purchase) is about 7900 MJ. Manufacturing is the most energy hungry phase, so that the author suggests concentrating efforts on reusing devices to extend their average usage time.

In detail the distribution of energy consumption [Moshnyaga 2010] (Figure 13) is as follows:

— Design/Manufacture: 4250 MJ (54%)
— Transport: about 950 MJ (12%)
— Purchase/Use: about 1500 MJ (19%)
— Other 1200 MJ (15%)

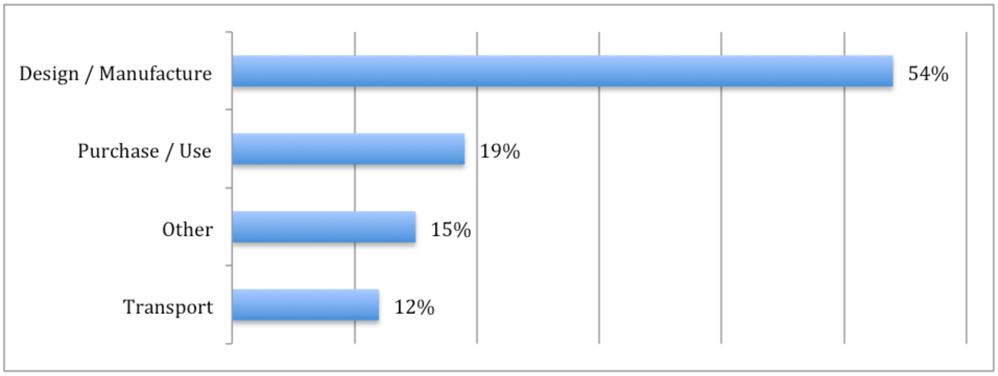

Fig. 13. Phases and Energy Costs

According to [Moshnyaga 2010] the energy taken for the production of a common Personal Computer has risen by 7% in comparison to 2002. This value increased from about 6420 MJ in 2002 to about 6900 MJ in 2007. About the usage phase, modern PCs consume more energy at full-load than the old ones, while in a low-power mode they take less energy than the old computers. Therefore the total energy used depends strongly on the usage scenario. Considering a home scenario, the total energy used by an average 2007 PC per year is almost the same as it was in 2002. The reduction is mainly due to CRT monitors (between 65 W-145 W when active, and 9-14 W in standby) substituted by LCD monitors (25 W when active, 2 W in standby). The increase is due to a possible increase in consumption of other components of a personal computer (graphics cards, memories, etc.) Thus, the overall (manufacturing + usage) energy consumption of PCs has increased over the last 10 years. According to the Wikipedia definition [9], despite it is a not official source, we can draw a trend about the frequency and TDP behaviour of AMD CPUs since 1996 as described in Figure 14.

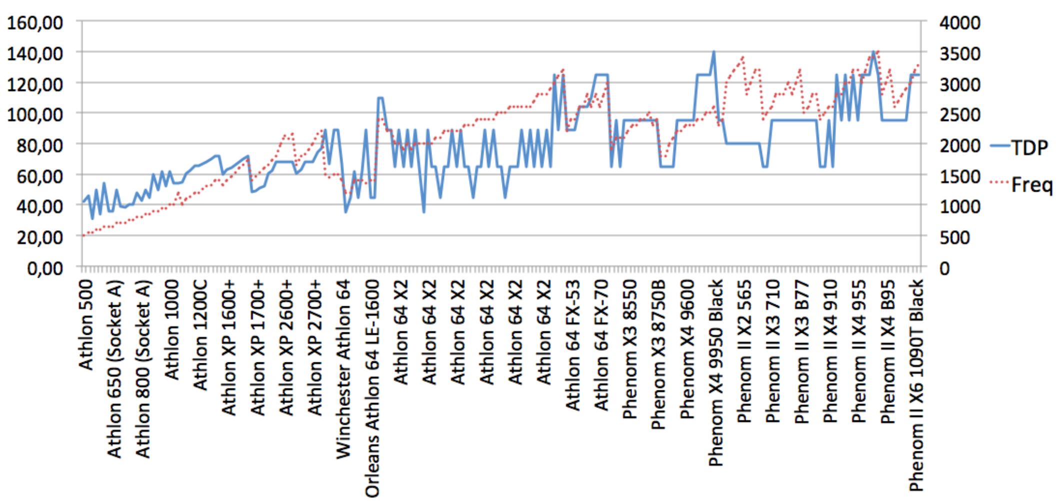

Fig. 14. CPU Frequency and TDP trend

[9] `http://en.wikipedia.org/wiki/List_of_CPU_power_dissipation` Last Visited: 23 Jan 2012

Without taking into account punctual values we can see that over this 15 years both frequency and TDP raised of respectively 7 and 3 times. We must point out that nowadays CPUs have multiple cores so TDP and Power Consumption are not the right metrics to measure CPUs energy efficiency.

4.4.2. Data Center, and Infrastructure. According to [Berl et al. 2009], the major challenge in energy reduction talking about "Cloud Computing" is the relation among system components and an optimal balance between performance, QoS, energy consumption, and self-aware runtime adaptation. Amazon [Hamilton 2009] calculated that the cost and operation of servers cost the 53% of their total budget, while energy-related costs reach 42% of the total. Figure 15 shows a typical monthly cost distribution in a data center.

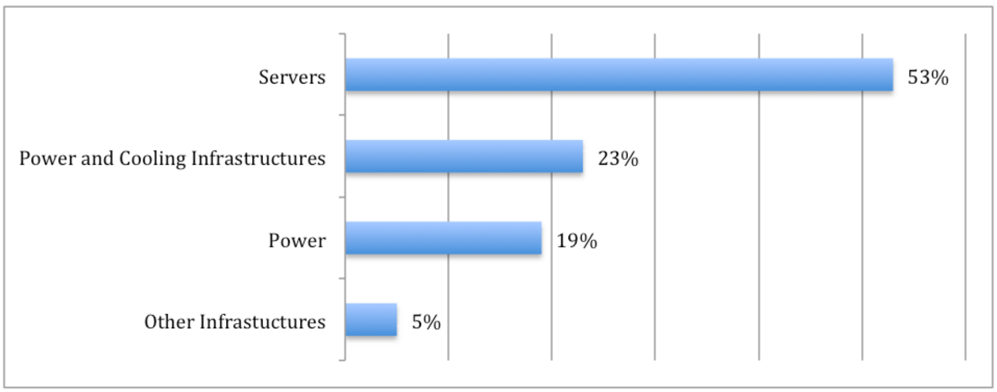

Fig. 15. Typical data center monthly costs distribution [Berl et al. 2009]

According to [Koomey], the electricity used in global data centers in 2010 accounted for between 1.1% and 1.5% of total electricity use. This means that data centers are using less energy than predicted by Environmental Protection Agency in 2007.

4.5. Considerations

As a conclusion, an analysis of the trends in IT energy consumption and emissions shows that the IT sector is responsible for a minority of them. However, the trend is increasing, due to the large increase in the number of individual IT devices (PCs and mobile phones), and the increase in energy used to manufacture them and, in using them. It is a responsibility of the IT sector to research ways to become sustainable, a further reducing emissions and consumption, especially at the level of individual devices.

5. ENERGY EFFICIENCY: GUIDELINES

As shown in Section 3, the literature proposes many tools to measure Energy Efficiency of IT devices. The next point is how to obtain Energy Efficiency, especially considering the OS and application layers and not only the hardware layer. We have surveyed the literature and we have found guidelines from Intel and Siemens [Kaefer 2009] [Steigerwald et al. 2007] [Larsson 2011] and also from academy [Gude and Lago 2011] We summarize in Table VII these guidelines according to the taxonomy proposed in Section 2.2.2. Each guideline could have an impact on the measures reported in section 3.

Table VII: Guidelines proposed by [Gude and Lago 2011] [Kaefer 2009] and [Larsson 2011]

ID	Layer	Guideline	Guideline Description
I-01	Infrastructure	Consider Cloud Platforms for energy efficient Internet applications [Kaefer 2009]	Cloud platforms use virtualization. This should improve energy efficiency of the internet application.
I-02	Infrastructure	Load balancing [Gude and Lago 2011]	Distributing workload evenly across two or more computers, network links, CPUs, hard drives, or other resources.
I-03	Infrastructure	Provide information for system management tools to support overall optimization [Kaefer 2009]	The use of power meters and energy aware applications provides information related the infrastructure energy consumption. This data should be used as an input to support the energy optimization.
A-01	Application	Efficient UI [Gude and Lago 2011]	Efficient UIs can let the user complete a task quickly. An inefficient UI can increase the application complexity, and consequently the energy consumption
A-02	Application	Decrease algorithmic complexity [Gude and Lago 2011] [Larsson 2011]	Reducing the algorithm complexity to accomplish a certain can save energy.
A-03	Application	Use Event-Based Programming [Gude and Lago 2011] [Kaefer 2009] [Larsson 2011]	Event based programming avoids a waste of resources involved in doing unnecessary operations. If polling cannot be avoided, it is advised to select a fair time interval.
A-04	Application	Use low-level programming and avoid use of byte-code [Gude and Lago 2011] [Larsson 2011]	With low-level programming languages, developers have more details of the system in which she/he is developing, than using high-level programming languages. When possible, it is advised to develop the more computationally intensive parts of the application in low-level Programming languages to increase performances and energy efficiency. The virtual machine interpretation of byte code can make the application energy inefficient.
A-05	Application	Put application to sleep [Gude and Lago 2011]	An application in sleep mode saves energy. An event, a signal, or an interrupt can resume the application.
A-06	Application	Batch I/O [Gude and Lago 2011] [Kaefer 2009] [Larsson 2011]	Buffering I/O operations increases energy efficiency; the OS can power down I/O devices when not used.

A-07	Application	Code Migration [Gude and Lago 2011]	To increase energy efficiency in devices where computation can be energy consuming, it may be worth moving the task to another energy efficient environment and gathering results when available.
A-08	Application	Reduce transparency/ abstractions [Gude and Lago 2011]	Layers of abstraction may cause inefficiency because no details from the underlying layer are exposed.
A-09	Application	Reduce data redundancy [Gude and Lago 2011] [Larsson 2011]	Storage and transportation of redundant data lowers energy efficiency
A-10	Application	Reduce QoS/Scale dynamically [Gude and Lago 2011] [Kaefer 2009] [Larsson 2011]	The application has to be able to change its behaviour in case of low-power situations.
A-11	Application	Reduce memory leaks [Gude and Lago 2011]	With memory leaks the application can stall or crash. This unpredictable behaviour can alter the energy consumption and, more generally, they must always be avoided.
A-12	Application	Use Power/energy profiling tools [Kaefer 2009]	This kind of applications model the energy behaviour of the system in which are run. This model should help the developer to optimize the application energy usage.
A-13	Application	Energy Efficiency code Patterns [Vetro' et al. 2011]	Static analysis techniques can highlight some code patterns that waste energy.
O-01	OS	Implement Power Management APIs [Kaefer 2009] [Larsson 2011]	Operating System can export power management APIs to enable applications to manage energy efficiency at lower levels.
O-02	OS	Optimal use peripherals [Gude and Lago 2011]	The use of a correct power management feature exploits properties and characteristic of peripherals
O-03	OS	Handle external signals and events [Gude and Lago 2011]	External signals and events: if incorrectly managed, can suspend the low power state of a (sub)system.
O-04	OS	Software-hardware interaction optimizations [Gude and Lago 2011]	A good design decision is to select which features are implemented by hardware, and which functionalities are developed by software. OS can take advantage of energy efficient hardware solutions with good drivers that make available these features to the SW layer.

O-05	OS	Use Compiler Optimization [Larsson 2011] [Gude and Lago 2011]	The Compiler can optimize the source code according to specific platform architecture.
O-06	OS	Use JIT Compiler [Kaefer 2009] [Gude and Lago 2011] [Larsson 2011]	A Just In Time compiler translates bytecode into machine languages at runtime; it can increase performance and decrease energy consumption.
O-07	OS	Use certified drivers for energy management (to allow idle states) [Kaefer 2009]	Drivers capable of exporting features to monitor and control power of the device support energy efficiency.
O-08	OS	Use only required services and background processes [Larsson 2011]	Unused applications waste memory, resource and energy.
H-01	Hardware	Power down peripherals [Gude and Lago 2011]	Peripherals can be storage, I/O, or network devices, GPS modules, etc. When not in use they should be set to a low power state or shut down.
H-02	Hardware	Lower the clock frequency [Gude and Lago 2011]	Reducing clock frequency causes less performance, less heating, and less energy consumption. QoS must be analysed before lowering clock frequency.
H-03	Hardware	Use specific-purpose hardware [Kaefer 2009]	A general-purpose hardware can be oversized for the specific problem. Oversized hardware can be translated in energy inefficiency.
H-04	Hardware	Dynamic Power Management Capabilities [Kaefer 2009]	Exploit features such as ACPI, processor and idle management, configurable high performance vs. low power components, to manage energy consumption.
H-05	Hardware	Power/ Energy metering support [Kaefer 2009]	Hardware metering support is often needed by profiling tools to get a reliable estimation of device power consumption.
N-01	Network	Efficient data traffic [Gude and Lago 2011]	Sending less data over the network can reduce energy consumption because the network interface is left in idle for more time. In the efficiency evaluation of this technique it is necessary to take into account the additional computation needed to implement data compression, proxying, etc.
N-02	Network	Tradeoff between compressed or raw data transfer [Kaefer 2009]	There can be an energy consumption overhead due to data compression but, after that, less data is exchanged.

N-03	Network	Energy impact of communication protocols [Kaefer 2009]	An energy efficient communication protocol can reduce the amount of information exchanged.

Most of the guidelines suggested in the literature are not strictly code-related, but they are mainly high-level recommendations for programmers and software designers (e.g. implement lazy loading of libraries). However, it is worth mentioning that such guidelines, despite being intuitive and acknowledged as effective by software industry specialists, did not receive any empirical validation. For this reason, an empirical validation that quantitatively assess their impact on Energy Efficiency is needed. Unfortunately it is not yet possible to provide evidence to demonstrate their effectiveness. Some of them such as I-02, A-02, A-03, A-09, A-13, O-05, and N-02 are more representable in empirical experiments, and may be assessed more easily than others.

6. CONCLUSIONS

Green IT is becoming a popular topic, but no specific surveys are available yet. In this paper we reported a survey of the literature about Energy Consumption & IT systems, starting from the viewpoint of Green IT. The survey has been performed searching for the following keywords: “Green IT”, “ICT Energy Consumption Reduction”, “Energy Efficiency”, “Energy Measurement”, “Power management”, “Energy Consumption Analysis” in the following databases: IeeeXplore, ACM digital library, IET Electronic Library and, more generally, Google Scholar. In order to organize the large number of papers found we have defined a taxonomy, based on two axes, the time axis (with activities such as design, manufacture, transport, use, dismiss and possibly recycle) and the space axis (with physical components of varying sizes, from larger to smaller: the clouds, data centers, computing nodes such as PCs, smartphones and mobile phones, applications, OS and hardware). First of all we have tried to contextualize the consumption of energy and resources of IT vs the rest. In 2007 IT electrical energy consumption in usage phase is reported to be 830 TWh or 0,5 % the total. In percentage this is a minimal amount, however in absolute terms it is relevant. Besides, estimates of IT consumption in the future show a fast growing trend. While we do believe that these figures are a good starting point, it should be noted that the accuracy of data reported is questionable. For sure consumption data is in many cases referred to several years ago (2007, 2009 for IT consumption) and their precision is not reported. At this regard a lot of work should be done to define and standardize the way consumption data is collected and reported. A first observation on the space axis is that there is no agreement on how to consider smart phones and mobile phones in general. Sometimes papers do not include them (strict IT and Green IT), sometimes they do (ICT and Green ICT), and sometimes the point remains fuzzy. Overall our point of view is that, considering the convergence of mobile phones into Internet nodes, they should be included. Besides, nowadays the production (and therefore the related consumption) of mobile phones is much larger than the one of computers. Most of the literature is about the space axis, and mainly about the usage phase. However, considering PCs at least, a study [Moshnyaga 2009] shows that the main contribution (about 50%) of energy consumption of a PC is due to the design/manufacture phase while the usage phase contribution represents only 20% of the total. The energy consumption of the usage phase, which currently does not reach 1% of the total, can be considered the one in which it is possible to intervene in a more distributed way. Small and distributed energy reductions can lead to large reductions worldwide. For this reason, an intervention on energy consumption reduction from a software point of view is to be considered interesting. We did not find similar studies on mobile phones or data centers. However, if the trend is confirmed, efforts to reduce energy consumption should concentrate on the manufacturing phase and/or on increasing the duration of the usage phase. Again considering the space axis, in 2009 data centers are the main users of energy (330 TWh) followed by (PCs, smartphones, tablets) (254 TWh) and network equipment (6TWh). From these figures is clear that data centers and PCs/smartphones should be the focus for energy consumption reduction during the usage phase. After this data collection and analysis phase, we have focused on methods and techniques to reduce energy consumption. At this regard we need precise ways to measure if consumption is actually reduced. So before all we have summarized the measures that can be used. Besides the obvious ones (energy and power) we have surveyed what has been proposed, and placed it into our taxonomy. Basically all measures proposed by different authors can be classified as measures of efficiency or productivity, applied to a node of the taxonomy in the usage phase. For instance efficiency for a data center is the ratio between energy for computation and total energy used (including conditioning). Finally, we have surveyed for techniques (or guidelines) to reduce consumption, and organized them in our taxonomy. What is available is a good starting point, but in many case the guidelines are quite high level, so their effect on consumption is hard to express in quantitative terms. In summary, future work by the Green IT community should be devoted to:

— Collect more precisely data about consumption, standardizing the data collection process;
— Collect and analyse in more depth consumption in non usage phase;
— Develop more extended and more detailed guidelines for energy consumption reduction;
— Validate and characterize quantitatively the effect of these guidelines.

REFERENCES

AMD. 2010. The truth about power consumption starts here. Tech. rep., AMD.

BERL, A., GELENBE, E., DI GIROLAMO, M., GIULIANI, G., DE MEER, H., DANG, M. Q., AND PENTIKOUSIS, K. 2009. Energy-efficient cloud computing. *The Computer Journal 53,* 7, 1045–1051.

BIANZINO, A. P., CHAUDET, C., ROSSI, D., AND ROUGIER, J.-L. 2010. A survey of green networking research. *CoRR abs/1010.3880*.

BIANZINO, A. P., RAJU, A. K., AND ROSSI, D. 2011. Apples-to-apples: a framework analysis for energy-efficiency in networks. *SIGMETRICS Perform. Eval. Rev. 38,* 3, 81–85.

CAPRA, E. AND MERLO, F. 2009. Green it: Everything starts from the software. In *ECIS*, S. Newell, E. A. Whitley, N. Pouloudi, J. Wareham, and L. Mathiassen, Eds. 62–73.

CEUPPENS, L., SARDELLA, A., AND KHARITONOV, D. 2008. Power saving strategies and technologies in network equipment opportunities and challenges, risk and rewards. In *Proceedings of the 2008 International Symposium on Applications and the Internet*. SAINT '08. IEEE Computer Society, Washington, DC, USA, 381–384.

FORGE, S. 2007. Powering down: remedies for unsustainable ICT. *Foresight 9*, 3–21.

GRID, T. G. Green grid metrics: Describing datacenter power efficiency.

GUDE, S. AND LAGO, P. 2010. a survey on green it: Metrics to express greenness in the it industry. Tech. rep., VU University.

GUDE, S. AND LAGO, P. 2011. Best practices for energy efficient software. *http://wiki.cs.vu.nl/green_software/index.php/Best_practices_for_energy_efficient_software*.

GUPTA, M. AND SINGH, S. 2003. Greening of the internet. In *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*. SIGCOMM '03. ACM, New York, NY, USA, 19–26.

GURUMURTHI, S., SIVASUBRAMANIAM, A., IRWIN, M. J., VIJAYKRISHNAN, N., KANDEMIR, M., LI, T., AND JOHN, L. K. 2002. Using complete machine simulation for software power estimation: The softwatt approach. *High-Performance Computer Architecture, International Symposium on 0*, 0141.

HAMILTON, J. 2009. Cooperative expendable micro-slice servers (cems): low cost, low power servers for internet-scale services. *4th Biennial Conference on Innovative Data Systems Research (CIDR) Asilomar*.

HYLICK, A., SOHAN, R., RICE, A., AND JONES, B. 2008. An analysis of hard drive energy consumption. In *MASCOTS*. 103–112.

IDC. 2009. International data corporation, document no. 221346.

JAIANTILAL, A., JIANG, Y., AND MISHRA, S. 2010. Modeling cpu energy consumption for energy efficient scheduling. In *Proceedings of the 1st Workshop on Green Computing*. GCM '10. ACM, New York, NY, USA, 10–15.

KAEFER, G. 2009. Green sw engineering: Ideas for including energy efficiency into your software projects. *Post Proceedings ICSE 2009*.

KANSAL, A., ZHAO, F., LIU, J., KOTHARI, N., AND BHATTACHARYA, A. A. 2010. Virtual machine power metering and provisioning. In *Proceedings of the 1st ACM symposium on Cloud computing*. SoCC '10. ACM, New York, NY, USA, 39–50.

KIRBY, A., BOGDANOVIĆ, J., AND PROGRAMME, U. N. E. 2008. *Kick the Habit: A Un Guide to Climate Neutrality*. UNEP.

KIYANCLAR, N. 2005. A survey of virtualization techniques focusing on secure on-demand cluster computing.

KOOMEY, J. Growth in data center electricity use 2005 to 2010. Tech. rep., Stanford University.

KRAVETS, R. AND KRISHNAN, P. 1998. Power management techniques for mobile communication. In *Proceedings of the 4th annual ACM/IEEE international conference on Mobile computing and networking*. MobiCom '98. ACM, New York, NY, USA, 157–168.

KRIKKE, J. 2008. Recycling e-waste: The sky is the limit. *IT Professional 10,* 1, 50–55.

LARSSON, P. 2011. Energy-efficient software guidelines. 14.

MAHADEVAN, P., SHARMA, P., BANERJEE, S., AND RANGANATHAN, P. 2009. A power benchmarking framework for network devices. In *Proceedings of the 8th International IFIP-TC 6 Networking Conference*. NETWORKING '09. Springer-Verlag, Berlin, Heidelberg, 795–808.

MILES, T 2010. Why is green important? impact of ict on the environment and energy use. Global Forum 2010.

MOLLA, A. 2009. Organizational motivations for green it: Exploring green it matrix and motivation models. In *PACIS*. 13.

MOSHNYAGA, V. G. 2009. Integrated circuit and system design. power and timing modeling, optimization and simulation. Springer-Verlag, Berlin, Heidelberg, Chapter Untraditional Approach to Computer Energy Reduction, 82–92.

MOSHNYAGA, V. G. 2010. A new approach for energy management in user-centric applications. In *Proceedings of the International Conference on Green Computing*. GREENCOMP '10. IEEE Computer Society, Washington, DC, USA, 107–112.

MURUGESAN, S. 2008. Harnessing green it: Principles and practices. *IT Professional 10,* 1, 24–33.

MURUGESAN, S. 2010. Making it green. *IT Professional 12,* 2, 4–5.

NAUMANN, S., DICK, M., KERN, E., AND JOHANN, T. 2011. The greensoft model: A reference model for green and sustainable software and its engineering. *Sustainable Computing Informatics and Systems 1,* 4, 294–304.

NORDMAN, B. AND CHRISTENSEN, K. 2005. Reducing the energy consumption of network devices. *IEEE 8023 tutorial*, 1–30.

PATEL, S 2011. Iea: Global power demand to surge 2.2% annually through 2035. http://www.powermag.com/renewables/waste_to_energy/3286.html.

RIVOIRE, S., SHAH, M. A., RANGANATHAN, P., AND KOZYRAKIS, C. 2007a. Joulesort: a balanced energy-efficiency benchmark. In *Proceedings of the 2007 ACM SIGMOD international conference on Management of data*. SIGMOD '07. ACM, New York, NY, USA, 365–376.

RIVOIRE, S., SHAH, M. A., RANGANATHAN, P., KOZYRAKIS, C., AND MEZA, J. 2007b. Models and metrics to enable energy-efficiency optimizations. *Computer 40*, 39–48.

ROTH, K. W., GOLDSTEIN, F., AND KLEINMAN, J. 2002. Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings Volume I: Energy Consumption Baseline. Tech. rep., US Department of Energy.

RUSSELL, J. T. AND JACOME, M. F. 1998. Software power estimation and optimization for high performance, 32-bit embedded processors. In *Proceedings of the International Conference on Computer Design*. ICCD '98. IEEE Computer Society, Washington, DC, USA, 328–.

SOMAVAT, P., JADHAV, S., AND NAMBOODIRI, V. 2010. Accounting for the energy consumption of personal computing including portable devices. In *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*. e-Energy '10. ACM, New York, NY, USA, 141–149.

STAUB, J. 2010. International energy outlook. *Outlook 0484,* July, 70–99.

STEIGERWALD, B., CHABUKSWAR, R., KRISHNAN, K., AND DE VEGA, J. 2007. Creating energy - efficient software. *Context 184,* 1, 30.

TIWARI, V. AND LEE, M. T.-C. 1995. Power analysis of a 32-bit embedded microcontroller. In *Proceedings of the 1995 Asia and South Pacific Design Automation Conference*. ASP-DAC '95. ACM, New York, NY, USA.

VETRO', A., TORCHIANO, M., AND MORISIO, M. 2011. Quantitative assessment of the impact of automatic static analysis issues on time efficiency. In *Informatica Quantitativa*.

WEBB, M. AND AL., E. 2008. Smart 2020: Enabling the low carbon economy in the information age. *The Climate Group London*.

WILLIAMS, E. 2004. Energy intensity of computer manufacturing: hybrid assessment combining process and economic input-output methods. *Environmental science technology 38,* 22, 6166–6174.