Evaluating Network-Based DoS Attacks Under the Energy Consumption Perspective

New security issues in the coming green ICT area

Francesco Palmieri
Dipartimento di Ingegneria
dell'Informazione
Seconda Università di Napoli
Aversa, Italy
francesco.palmieri@unina.it

Sergio Ricciardi
Departament d'Arquitectura
de Computadors
Universitat Politècnica de Catalunya
Barcelona, Spain
sergior@ac.upc.edu

Ugo Fiore
Centro Servizi Informativi
Università degli Studi
di Napoli Federico II
Naples, Italy
ufiore@unina.it

Abstract *— In the green Information and Communication Society (ICS), new form of Denial of Service (DoS) attacks may be put in place: exploiting the computational and storage resources of datacenters with the aim of consuming as much energy as possible, causing detrimental effects, from high costs in the energy bill, to penalization for exceeding the agreed quantity of CO_2 emissions, up to complete denial of service due to power outages. To the best of our knowledge, this is the first paper which investigates the impacts of network-based DoS attacks under the energy consumption perspective. We analyzed different types of such attacks with their impacts on the energy consumption, and showed that current energy-aware technologies may provide attackers with great opportunities for raising the target facility energy consumption and consequently its green house gases (GHG) emissions and costs.*

Keywords: denial of services DoS; energy consumption; networking; datacenters; green house gasses emissions GHG.

I. Introduction

In the last years, the emergence of the new green, energy-sustainable computing paradigms has gained a lot of attention in both the research and industrial arenas. Consequently, the development of modern ICT architectures with the additional requirement of keeping the energy consumption under control while maintaining the services offered at a satisfactory level, according to the new concepts of energy-efficiency and energy-awareness, has become central. By observing the electrical power demands of the largest networked computing farms, such those empowering the modern distributed computing and cloud infrastructures, it has been estimated [1][2] that ICT worldwide energy consumption amounts to more than 8% of the global electricity production and the energy requirements of datacenters, storage and network equipment are foreseen to grow by 12% per year. Clearly, such a huge electricity demand will result in environmental and engineering issues and bottlenecks, seriously conditioning the evolution of the whole ICT sector. For example, we can consider that the number of transistors integrated within the recent Intel Itanium processors reaches to nearly 1 billion of elements. If this growth rate continues, the heat (per square centimeter) produced by next-generation CPUs would exceed that of the sun's surface [3], by reaching a critical technological limit and energy demand threshold. Furthermore, together with the growth of the energy required by the above infrastructures, there is an alarming rise in their correct usage involving thousands of concurrent e-commerce transactions and millions of Web queries per day, handled through large-scale distributed datacenters, which consolidate hundreds or thousands of servers with other auxiliary systems such as cooling, storage and network communication ones. In this scenario, there has been an equally dramatic evolution in security. The need for efficient ways of detecting and attempting to prevent intrusions, as well as of mitigating attacks, has led to the elaboration of sophisticated analysis techniques and countermeasures. This has brought a corresponding advance in the cleverness of attack strategies and tools, also affecting the attack objectives that can become different from the traditional ones (confidentiality, integrity, availability or performance of the computing elements offering the service). The developments in the areas of energy-awareness/efficiency and network/site security have been considerable but separate. This paper underlines that there are areas in common between these two fields, and addresses a new perspective, which might become commonplace over the next years: attacks could change in their main aims, either exploiting weaknesses in power-saving and management mechanisms to disrupt services, or even attempting to increase the energy consumption of an entire farm, by causing financial damages. Therefore, it becomes clear that the energy-efficiency and security challenges can be better addressed in a combined way if the energy requirements and the bottlenecks of the underlying security technologies and protocols are better understood and coped accordingly. For this sake, we evaluated the efficiency of common attacks with respect to their troublemaking potential in terms of the impact on the energy consumption of the target infrastructure. To the best of our knowledge, this is the first paper that evaluates the DoS attacks under the energy consumption perspective.

II. Energy-Oriented Denial Of Service Attacks

Denial of Service attacks are becoming a more and more important disturbance factor on all the sites that are connected to the Internet. Any defense against these menaces is very difficult because they strive at consuming all the

978-0-7695-4532-5/11 $26.00 © 2011 IEEE
DOI 10.1109/BWCCA.2011.66

available resources at both the computing service and network transport layers, where it is very hard to distinguish whether an access or service request is genuine or malicious. By affecting the server systems or the network connection on the target sites, the attacker may be able to prevent any access to e-mail relays, websites, online accounts (banking, e-commerce, etc.) or other services that rely on them [4]. Attacks to the network connection take place by exhausting the available bandwidth through the generation of a very a large number of packets directed to the target site. Typically, these packets are ICMP ECHO packets but in principle they may be anything [5]. On the other hand, the computing resources on the service nodes within the target site can be saturated by overwhelming them with a huge quantity of CPU intensive service requests, such repeated transaction attempts on an HTTPS or any kind of SSL-empowered server. In order to increase the attack power, many remotely controlled computers can be simultaneously used as the source. This kind of menace is also widely known as Distributed Denial of Service (DDoS) attack. We advise the possibility that the above menaces can be made more effective, by introducing new and more subtle objectives and attack scenarios, based on power management and energy-efficiency considerations. A fundamental property of such kind of attacks is that they exploit the hardware components and subsystems that experience the heaviest difference in power consumption between their busy and idle/sleeping states. The attacks perform their offending activity by keeping the target component as busy as possible, and preventing it from going into low power usage modes, and thus forcing it to work at its near-maximum speed, frequency, voltage or temperature. In modern ICT infrastructures, the most critical components from the power consumption perspective are the server systems, whose energy demand is tightly related to their load: a fully loaded server absorbs about two times the power of an idle one, with a linear increment of the power consumption with respect to the server load. The CPU alone contribution to the server power consumption goes from 25% to 55%, depending on the server, followed by memory and networks interfaces [6][7]; disks, motherboard and fans consume less energy (Table 1). Anyway, it should be pointed out that, in order to assess the potential of an energy-oriented attack, it is not a priority to focus on the major power hungry device, but rather on the most energy sensible devices, i.e. components whose energy consumption strongly varies with the traffic load. These components are mainly CPU, disks, network interface cards (NIC); in the Section III we focus on such components and analyze their impact on the energy consumption. The first and most critical component that exhibits these operating characteristics is the CPU/Memory subsystem whose energy consumption is known to scale linearly with its utilization [6][8]. Since the goal of such energy-oriented attacks is to maximize the power consumption by keeping the CPU and memory on the target systems as busy as possible, they try to add additional load on the servers by introducing a large number of service request which subtract most of the resources to the legitimate ones and let the CPUs working at their maximum operating frequency. This can be achieved by overwhelming the CPU/Memory runtime subsystem with fake SSH or SSL/TLS-based transactions and service requests or forcing the continuous execution of a huge number of random read and write operations on very large arrays located in memory to generate a large quantity of cache misses.

TABLE I. ENERGY CONSUMPTION BREAKDOWN OF A LOW-END SERVER

Component	*Peak Power*
CPU [19]	80 W
Memory [20]	36 W
Disk subsystem [21]	12 W
Network Interface [22]	2 W
Motherboard [6]	25 W
Fans [6]	10 W

CPU and memory dominate power absorption; disks are relevant only if there are many.

Another effective way of draining more and more system energy is overloading the device's hard disks with millions of read or write operations by forcing them to constantly operate at their maximum sustained transfer rate or to continuously spin up and down the hard disks spindle engines. This kind of attack is very common in the offending strategies of several computer viruses and Trojans that are typically able to directly run malicious codes on the target nodes. In the worst cases, the malicious agents can alter the operating system kernel or some application binary code so that more energy is needed for their execution. However, the binaries altered in such a way may or not continue to behave correctly from the users' point of view. Finally, the last device/component that can be solicited is the network interface, when its energy consumption depends on the actual connection rate, that is, in implementations supporting adaptive link rate technologies (ALR), low power idle (LPI) and dynamic voltage scaling (DVS) mechanisms. Clearly, the disruption of these attack schemes is dependent on how much power the device consumes in maximum speed mode respect to the one required in lower power modes; such a gap may be as high as 90% between idle and full load states for higher speed interfaces [9]. Perhaps unexpectedly, security systems themselves also offer a wealth of opportunities for energy-oriented attacks. Firstly, it must be remarked that security systems, while being essential to the correct operations of networked systems, also have an impact on the power expenditure. Such security systems strive to monitor the behavior of the device under control as non-obtrusively as possible. However, they consume a not negligible amount of energy [10]. For example, keeping in mind that, (a) in reasonably well-managed organizations, end-user PCs are ordinarily equipped with properly configured and updated antivirus software, and (b) this software will scan on-the-fly some or all the content trying to reach the computer local storage, and (c) there are conditions under which the antiviral scanning causes long periods of full CPU load, a disruptive energy-aware attack can be orchestrated in the following way. Firstly, the attacker selects a (ideally innocent) content, which will trigger the antivirus reaction, consuming a great amount of CPU in the process. Secondly, the attacker sets up a Web site with a content appealing to the individuals belonging to the target organization (alternatively, he/she may set up a spam campaign) and have the malicious content delivered to the target. Spam, that is a nuisance to email users, since it eats up remarkable resources spent to deal with it and prevents it from reaching user mailboxes, can also be exploited for energy-oriented attacks. Spam messages are usually cheap to send, because they are normally originated from compromised computers belonging to a botnet. Mail servers run anti-spam software, which has the purpose of identifying and filtering out unwanted messages, and this software consumes CPU, disk and network resources [11]. An energy-oriented attack could then increase the footprint on a target mail server by simply increasing the amount of spam addressed to it. In any case, a successful attack will maximize power consumption and excessively solicit hardware components while presenting to the user the appearance that the system is behaving normally, with the possible exception of an increased CPU, disk or network activity. Some of the side effects that one would expect to observe in presence of these menaces, if they are not implemented more subtly, include the legitimate user requests being served slowly, the CPU fan turning on while the user is performing some action that does not normally cause the fan to come on, the system becoming less interactive than usual, the network loosing part of its speed/responsiveness and the hard drive spinning up immediately after a spin down.

A. Affecting the energy costs

Incrementing the power usage has direct and immediate consequences on the energy expenses. If well designed, attacks may exploit the different energy costs (e.g. during the the night-day cycle) or the energy budget threshold that the facility agreed with the power supplier, resulting in very high energy bills. That's worse, traditional power provisioning strategies, aiming at keeping as much computing and storage equipment as possible within a given power budget in order to maximize the utilization of the deployed datacenter power capacity, may present the drawback of offering more subtle vulnerabilities to possible attackers. More precisely, such strategies try to fill the gap between achieved and theoretical peak power usage in order to deploy additional equipment within the power budget [6]; the full utilization of the datacenter is offset by the risk of exceeding its maximum capacity, resulting in power outages or costly SLA violations due to the fact that the maximum drained power of a datacenter may be conditioned by a physical and/or contractual limit. The contractual enforcement exceeding will result in economic penalties (that can be exploited by a malicious competitor), or even overcoming the physical power limits resulting in power outages.

B. Neutralizing energy saving systems

If attackers know that some energy-saving mechanisms operate in the target system/network, and if they know the details about these system, they can devise attacks aimed at neutralizing them. This is a subtle issue, because the amount of extra work to be "injected" into the system does not need to bring the processor or storage to full load, but is limited to the amount necessary to avoid the triggering of the energy-saving mechanisms, which are, in general, threshold-based. This means that detecting such attacks can be significantly harder. Furthermore, energy saving techniques are the more vulnerable to energy-oriented attacks, since they offer to the attackers greater opportunities to rise the energy consumption. It is quite common, in fact, that an infrastructure, like a datacenter, buys a given amount of energy to be used into an agreed period of time, according to the mean energy consumption of the site; exceeding such threshold may result in additional costs. An attacker may exploit such situation by raising the computational needs of the site and, thus, its energy consumption, above the threshold, therefore causing an economical damage or, even worse, an energy outage resulting in a complete denial of service. A very simple example of the above concepts can be observed when per-server sleep mode is deployed in the datacenters. An attack that simply generates continuous fake demands/traffic for all the servers may prevent machines to go into sleep mode during low load periods, thus having large impacts on the medium and long-term power consumption and hence conditioning the overall energy containment strategy.

C. Incrementing the operating temperature

Even if harder to put into practice, since it requires attacker to gain access to computing resources, thermal-based attacks are another potential menace that has to be taken into account. Such offensive strategies aim at executing a particular piece of code whose objective is not to saturate the computing or storage resources, but instead to subtly execute a relatively small cycle-loop that heats the CPU and the memory banks. The current CMOS technology provides modern microprocessors with not only transistors but also capacitors and resistors. Under normal circumstances, the CPU is not always active at 100%, but instead enters and exits from low power periods (HLT machine code instruction) in which the clock is halted and the circuitry enters a suspend mode until an interrupt or reset happens. Also, low power states (C-states) are available in the latest Intel® CPUs. Malicious codes may prevent the CPU to enter such low power states and continuously executes loops that charge resistors, notably increasing the temperature. Current datacenter infrastructures, in fact, have a power usage effectiveness (PUE) of 2, meaning that the heat, ventilation and air conditioning systems (HVAC) consume as much energy as the computing and storage resources. Therefore, the quantity of power absorbed by the HVAC system is not negligible: the potential of thermal-based attacks is as high as the energy-oriented one. The result is that the cooling infrastructure will work harder consuming a considerably higher quantity of energy. Detrimental effects of such attacks include the increase of the CPU and memory temperatures, with the consequent stability problems, reduced component life (an increase by 10°C halves the chip life span), and increased cooling power consumption.

D. *Exhausting the power budget*

As we have seen, attacks may result not only in high energy costs, but, in the worst cases, also in complete power outages. It has been observed that the power consumption declared by manufacturers (nameplate value) is usually an overestimated conservative value [12], and thus it is of limited usefulness when predicting the total power budget of the datacenter, giving the idea that "there will be enough power" if the nameplate values are considered when dimensioning the power facilities. This scenario, together with the periodical updates of new components (additional memory banks, disks, network interface cards, etc.) in an effort to accommodate the growing business demand and the wear of the devices as long as the substitution of old components with newer ones, which are more efficient but also more power-hungry (Moore's law has not been compensated at the same pace by energy efficiency), exposes the datacenter facility to the risk of exceeding its maximum power budget, in particular under energy-oriented attacks. Accordingly, we point out the security related risks of over-subscribing the datacenter under the energy consumption perspective. In fact, a sustained energy-oriented attack may put an entire datacenter out of service by totally blocking the underlying electrical distribution system (by exhausting its capacity). Such kind of attacks may be hard to detect, unless a constant fine-grained on-line monitoring and data-collection systems are deployed directly on the power distribution sub-system (i.e. UPS, PDU, RACKS, etc).

E. *Incrementing dirty emissions*

Energy-oriented attacks may also be exploited under an additional dimension: the green house gases emissions (GHG). Several practices have been adopted by the industry and the governments to reduce the GHG emissions [13]: carbon taxes, cap & trade, and carbon offset are all susceptible of being exploited by attackers to increase the GHG emissions of a facility and thus its costs. In a carbon tax approach, industries pay taxes according to the amount of emitted GHG (mainly CO_2); in this context, an attacker may obtain a double objective: raising both the energy consumption and the costs associated with the increased GHG emissions. In cap & trade containment strategy, a limit (cap) is imposed on the maximum allowed emissions and a market (trade) is created in which additional emission permissions may be bought by virtuous industries that do not reach the cap. In the carbon offset approach, industries are committed to compensate their emissions by buying in "green", such as tree reforestation, etc. Both the cap & trade and the carbon offset policies may attract unsavory practices from organizations that take advantages of third party emissions induced by the aforementioned attacks.

III. Modeling Power consumption in Energy-oriented DoSes

To illustrate the potential of energy-oriented attacks and analyze their dynamics and behaviors, we modeled the additional power consumption associated to each one of them. When exploiting the CPU/Memory subsystem, we consider that a modern CPU dynamically adapts its operating frequency to the current load so that its instantaneous power demand at the frequency *f* can be estimated as:

$$P(f) = \frac{1}{2} C V_f^2 A f \,. \tag{1}$$

In the above theoretical formulation [8], *f* can assume values within the range [f_{min}, f_{max}], *C* (aggregated load capacity) and *A* (activity factor) are fixed constant parameters (depending on the involved CPU characteristics), and V_f is the CPU voltage scaling linearly with the frequency *f*, that is:

$$V_f = V_{\max} \frac{f}{f_{\max}} , \tag{2}$$

where V_{max} is the maximum operating voltage required at the frequency f_{max}. Since the goal of all the CPU-based attacks is overloading the CPU by forcing it to work at its maximum operating frequency f_{max} for the longest possible time, we can estimate the worst case and best case power demands P_{max} and P_{min} as:

$$P_{\max} = \frac{1}{2} C V_{\max}^2 A f_{\max}, \quad P_{\min} = \frac{1}{2} C \left(V_{\max} \frac{f_{\min}}{f_{\max}} \right)^2 A f_{\min} \cdot \tag{3}$$

Consequently, if we consider that the average server utilization of datacenters is very low, often below 30% of its CPU capacity [6][14][15], we can assume that the average CPU power consumption approximates to P_{min} and hence the additional energy consumption introduced by a CPU based DoS attack can be estimated as:

$$E_C = (P_{\max} - P_{\min}) t_d = \frac{1}{2} C V_{\max}^2 A \left(\frac{f_{\max}^3 - f_{\min}^3}{f_{\max}^2} \right) t_d , \tag{4}$$

where t_d is the duration of the attack. Thus, the energy increase is proportional to the difference of the cubes of the maximum and minimum frequencies, and depends only linearly on the attack duration. This means that attack intensity is more critical than attack duration. Many bursty, strong attacks can achieve the same objective as one single sustained attack with lower intensity and, while the latter may be harder to detect but easier to prevent, the former will be easier to detect but harder to prevent. Analogously, the additional energy demand E_D for a typical attack based on repeated disk operations can be calculated by referring to the involved transfer rate *r* and considering the maximum sustainable drive transfer rate r_{max} as a worst case metric to calculate the amount of time spent in read mode when transferring data at speed *r*. We focus on attacks based on read operations since large-block-size reads consume more energy than writes (approximately P_{read}=13.3$\mu W/Kbyte$ against P_{write}=6.67$\mu W/Kbyte$ [16]), and the fact that reads occur 4-5 times more than writes becomes significantly important when considering that read operations may be used much more easily also on a partially compromised host. Let P_D be the power required by a disk (read) operation, as sum of engine-dependent mechanical power consumption [17], with the operation-dependent (read-write) electronic power consumption:

$$P_D = \frac{K^2\omega^2}{R} + D_r w_r P_{read}\,, \tag{5}$$

where K is a motor voltage constant, R is the motor resistance, D_r is the (read) demand (*kByte*) and w_r is a weighting factor depending on the current rate r [18]:

$$w_r = \frac{r}{r_{\max}}\,. \tag{6}$$

The first factor of eq. (5), referring to mechanical movement, quadratically depends on the angular velocity ω, and is the most critical part from the energy consumption perspective: an attack generating randomly sparse and bursty block read operations, forcing the disk hardware to continuously spinning up immediately after a spin down, can introduce a near maximum burden to the overall disk energy demand. Then, by considering that the energy required by the drive in low power P_l is already included in the default system power consumption, we can argue that the additional energy required during a disk-based DoS attack is upper-bounded by the above activities that can be expressed by:

$$E_D = (P_r - P_l)\cdot t_d\,, \tag{7}$$

where t_d is the duration of the attack. Note that eq. (7) is a function of the involved transfer rate r, so that the higher the sustained transfer rate during the attack is, the greater the impact on the overall power consumption will be. Finally, also if a more limited quantity of energy is required for the network interface, in presence of modern NICs supporting dynamic link rate adaptation or low power idle mechanisms (P_{min}), and hence reducing their speed and energy requirements in case of limited or no traffic, significant increments in power usage can be achieved by forcing the interfaces to work at their maximum throughput by flooding the target hosts with typical DDoS-generated traffic. Also in this case, if P_{max} is the power demand in active/maximum speed mode, the additional energy absorption introduced by an attack of duration t_d can be expressed by:

$$E_N = (P_{\max} - P_{\text{avg}})\cdot t_d\,, \tag{8}$$

where P_{avg} is the average power consumption of the NICs with normal traffic load during the time interval t_d.

IV. TESTS RESULTS AND DISCUSSION

Using the power consumption model of Section III, we estimated the potential of the CPU-based attack of eq. (4), for the reference processor AMD® Athlon® 2,4 GHz. Following the approach used in [8], we imposed a lower bound for $f_{min} = f_{max} / 2.4$ to prevent asymptotic trend during low utilization periods, where the constant 2.4 and the V_{max} = 1.4 V values are based on the frequency range and maximum voltage provided by the specification sheet of the reference CPU. As we can see from Fig. 1, the CPU-bound has great potentiality to exploit the energy consumption with a maximum intensity attack. The energy consumption surplus reached during the maximum intensity is up to 13.8 times the minimum energy consumption under low utilization periods and up to 4.1 times the energy consumption with medium load (absolute values scaled by constant factor $K = AC$).

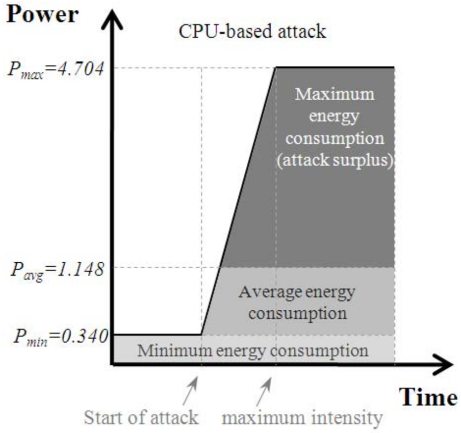

Figure 1. Power consumption upper bound for aCPU-based attack.

Power consumptions of the I/O-based attack are plotted in Fig. 2. The maximum read rate has been assumed in all the cases (i.e., r=r_{max}) and the variation of the power consumption has been reported for different values of the angular velocity ω∈{5400, 7200, 10000, 15000} *rpm* (absolute values scaled by constant factors K_1=K^2/R, K_2=D_rP_{read}). As we can see, the angular velocity strongly influences the disk power consumption, with the highest intensity attack that may reach peaks of 7.6 times the low power consumption mode.

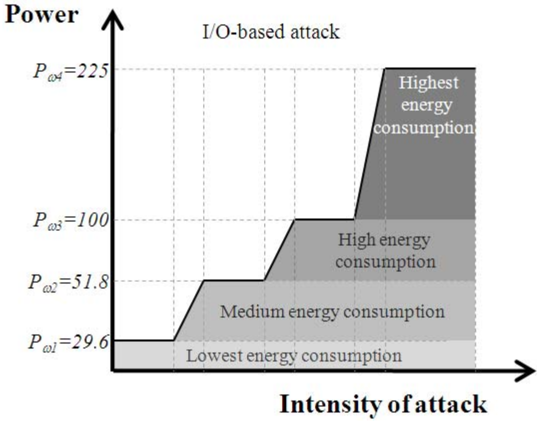

Figure 2. Power consumption upper bond for anI/O-based attack.

The CPU-bound based attack achieves the higher power consumption, while the I/O-bound one is less sensible to the power consumption, even if the latter may slow down the datacenter responsiveness even more than the former. The energy consumption potentiality of an energy-oriented attack on the NICs can be found in [23]. Nevertheless, the potential of the attacks should be contrasted also with the difficulties of being deployed. CPU and I/O-bound attacks are much easier to commit, since a DDOS may easily generate a huge number of web searches or mail requests, whilst the offline job processing requires access to back office facilities, like in the service as a service (SaaS) cloud computing paradigm.

V. Conclusions

DDoS attacks have the potential not only of denying the service of the target facility, but may be carved to explicitly impact its energy consumption. Such attacks may be targeted at several objectives: increment the energy consumption, the GHG emissions and introducing, in the worst cases, power outages. Some of these attacks are relatively easy to perform, e.g. CPU and I/O-bound based ones, whilst others are more difficult to deploy. In any case, the potential of such attacks should not be underestimated. Effective power management techniques have to be deployed to prevent detrimental effects. The most effective technique is the power capping scheme that set a maximum power consumption threshold and operate the facility always below that value. A power monitoring system constantly monitors the power absorption, and if an increment is detected, takes the corresponding actions to decrease the power, from job de-scheduling/ migrating to using any available component-level strategy to decrease the energy consumption, e.g. CPU voltage/ frequency scaling (DVS), downclocking devices, forcing sleep mode, etc., i.e. implementing an energy proportional computing system which has proved to be an effective way to reduce peak power usage. Anyway, it should be pointed out that power capping alone, although is an immediate measure to prevent facility detrimental, is not enough to detect attacks. Network based DoS attacks have to be recognized and isolated from the allowed traffic through a comprehensive security system.

Acknowledgment

This work was supported in part by the COST Action IC0804 on Energy Efficiency in Large Scale Distributed Systems, the Spanish Ministry of Science and Innovation under the DOMINO project (TEC2010-18522), the Catalan Government under the contract SGR 1140 and the DIUE/ESF under the grant FI-201000740.

References

[1] BONE project, 2009, “WP 21 Topical Project Green Optical Networks: Report on year 1 and updated plan for activities”, NoE , FP7-ICT-2007-1 216863 BONE project, Dec. 2009.

[2] J. G. Koomey, “Estimating total power consumption by servers in the U.S. and the world”, Lawrence Berkeley National Laboratory, Stanford University, 2007.

[3] G. Koch, “Discovering multi-core: Extending the benefits of Moore’s law”, Technology@Intel Magazine, 2005, http://www.intel.com/technology/magazine/computing/multi-core-0705.pdf.

[4] M. McDowell, “Understanding Denial-of-Service Attacks”, National Cyber Alert System, Cyber Security Tip ST04-015.2004, 2004.

[5] Denial of Service Attacks, 1999. Online, [2010.11.10] http://www.cert.org/tech_tips/denial_of_service .html.

[6] X. Fan, X-D. Weber, L.A. Barroso, “Power provisioning for a warehouse-sized computer”, in Proc. 34th annual international symposium on computer architecture (ISCA ’07), pp 13–23, 2007.

[7] L.A. Barroso, U. Hölzle, “The Case for Energy-Proportional Computing", IEEE Computer, vol. 40, pp. 33-37, 2007.

[8] D. Meisner, B.T. Gold, T.F., Wenisch, “PowerNap: eliminating server idle power”, In Proc. of ASPLOS ’09, pp 205–216, 2009.

[9] K. Christensen, P. Reviriego, B. Nordman, M. Bennett, M. Mostowfi, J. A. Maestro, “IEEE 802.3az: The Road to Energy Efficient Ethernet”, *IEEE Comm. Magazine*, Nov. 2010.

[10] J. Bickford, H. A. Lagar-Cavilla, A. Varshavsky, V. Ganapathy, L. Iftode, “Security versus Energy Tradeoffs in Host-Based Mobile Malware Detection”, MobySis 11, Bethesda, Maryland, USA, 2011.

[11] McAfee and ICF International, “The Carbon Footprint of Email Spam Report”, 2009.

[12] J. Mitchell-Jackson, J. G. Koomey, B. Nordman, and M. Blazek, “Data center power requirements: measurements from silicon valley”, Energy ISSN 0360-5442, 837–850, 2003.

[13] B. St Arnaud, “ICT and Global Warming: Opportunities for Innovation and Economic Growth”, http://docs.google.com/Doc?id=dgbgjrct_2767dxpbdvcf.

[14] C. Bash and G. Forman, “Cool job allocation: Measuring the power savings of placing jobs at cooling-efficient locations in the data center,” in Proc. of the 2007 USENIX Annual Technical Conference, 2007.

[15] P. Bohrer, E. Elnozahy, T. Keller, M. Kistler, C. Lefurgy, and R. Rajamony, “The case for power management in web servers,” Power Aware Computing, 2002.

[16] A. Lewis, S. Ghosh, and N.-F. Tzeng, “Run-time energy consumption estimation based on workload in server systems. In HotPower”, 2008.

[17] S. Gurumurthi, A. Sivasubramaniam, M. Kandemir, and H. Franke, “Reducing Disk Power Consumption in Servers with DRPM”, Computer 36, 12, pp. 59-66, 2003.

[18] W. West, E. Agu, “Experimental Evaluation of Energy-Based Denial-of Service Attacks in Wireless Networks”, IJCSNS International Journal of Computer Science and Network Security, VOL.7 No.6, Jun. 2007.L.A. Barroso, U. Hölzle,, “The Case for Energy-Proportional Computing", IEEE Computer, vol. 40, pp. 33-37, 2007.

[19] Intel Corporation, “Intel Xeon Processor with 512KB L2 Cache at 1.80 GHz to 3 GHz Datasheet”, http://download.intel.com/design/Xeon/datashts/29864206.pdf, Mar. 2003.

[20] Micron Technology Inc., “Calculating Memory System Power for DDR”, http://download.micron.com/pdf/technotes/ddr/TN4603.pdf, 2001.

[21] Seagate Technology LLC. Product manual Barracuda 7200.7. http://www.seagate.com/support/disc/manuals/ata/cuda7200pm.pdf, Sep. 2005.

[22] R. Sohan, A. Rice, A. W. Moore, K. Mansley, “Characterizing 10 Gbps network interface energy consumption”, LCN 2010, 268-271, 2010.

[23] S. Ricciardi, D. Careglio, U. Fiore, F. Palmieri, G. Santos-Boada, J. Solé-Pareta, “Analyzing Local Strategies for Energy-Efficient Networking”, in Proc. of SUNSET 2011, IFIP NETWORKING, Valencia, Spain, LNCS 6827, pp. 291-300, 2011.