Reducing Lifecycle Energy Use of Network Switches

Priya Mahadevan, Amip Shah, *Member IEEE*, Cullen Bash

***Abstract*—The networking infrastructure is becoming a growing portion of the energy footprint in data centers as well as information technology in general. In this paper, we evaluate energy management strategies for network switches. We begin by performing a lifecycle assessment of existing switches in a data center, and find that the use phase of the lifecycle dominates. We then parametrically examine various energy management techniques to reduce this operational footprint, and find that advances in operational energy efficiency may soon increase the share of environmental burden from switch manufacturing. We conclude by discussing how these findings may influence network design in the future.**

***Index Terms*—Computer networks, computer network management, environmental factors, life cycle costing**

I. Problem Addressed

The environmental footprint of information and communications technology (ICT) is rapidly becoming a matter of growing concern. It has been estimated that ICT is responsible for about 2% of global carbon emissions [1], roughly as much as the aviation industry. While numerous researchers have previously explored the lifecycle environmental impacts of various ICT systems, the role of the underlying network that connects different systems is a relatively less studied field. Particularly within data centers—which constitute one of the fastest growing portions of the ICT footprint [2]—the lifecycle impact of network switches is an understudied area.

Networking devices in data centers in the U.S. alone accounted for 3 billion kWh per year in 2006 and this number is growing rapidly [27]. However, reducing network power consumption has attracted attention only recently in the community. A first step in this mostly unexplored space is to understand the current trends in lifecycle energy consumption of data center networking devices such as switches. Based on this data, we further analyze and predict how aggressive power management schemes with a view to reducing switch operational energy impact their lifecycle energy.

Priya Mahadevan is with the Multimedia Communications and Networking Lab at Hewlett Packard Laboratories in Palo Alto, CA 94304 USA (e-mail: priya.mahadevan@hp.com).

Amip Shah and Cullen Bash are with the Sustainable IT Ecosystem Lab at Hewlett Packard Laboratories in Palo Alto, CA 94304 USA (e-mail: amip.shah@hp.com; cullen.bash@hp.com).

II. Knowledge of Prior Work

Numerous studies have explored lifecycle impacts of ICT systems such as mobile telephony [3], personal computers [4-7], and enterprise computing [8, 24]. The lifecycle energy use of network devices has also been studied [9, 10], but mostly in the context of home or office environments. Some specific studies have attempted to quantify energy efficiency in networks as function of data transfer rates [11], but these are generally empirical models at an industry level and cannot be mapped to specific switch design or operating characteristics.

Within enterprise and data center computing environments, several researchers have explored operational power measurement and management techniques for servers [12-15] and the facilities infrastructure [16-18]. In contrast, energy consumed by networking devices has been only recently targeted [9, 19-20, 29-30], mainly because they have historically taken up a relatively small portion of a data center's power budget [26]. As high-speed network connections become more widespread in home, enterprise and data center environments, and as they experience elevated utilization levels, we expect switches and routers to become a more significant fraction of the total data center footprint. Moreover, lifecycle considerations in data center networking infrastructures have been ignored to date. We fill this gap by presenting the first end-to-end approach for network energy management that includes both manufacturing and operational costs based on real production data center workload.

III. Project Undertaken

We propose an end-to-end analysis of a data center network. To determine the appropriate scope for our work, we consider current data center network architectures, where switch elements remain powered on at all times, irrespective of the traffic flowing through them [20]. We also consider the impact of a recently proposed network topology that aims to minimize the data center network energy consumption. Further, we also consider the impact of network power management schemes recently proposed in the literature.

Today's data center network topology typically consists of a two- or three-level tree of switching elements with progressively more specialized and power hungry equipment moving up the network hierarchy [21] [Fig. 1(a)]. The rack switches essentially operate on the 'edge' of the network, capturing user requests for data transmission. The aggregation switches accumulate the traffic from a variety of rack switches and transmit them towards a central destination. The core switches essentially provide the interface between a data center and the outside world.

By contrast, a recent study [21] experimented with a special instance of an alternative topology (called a "fat-tree") to interconnect commodity Ethernet switch silicon [Fig. 1(b)]. While the fat-tree employs more individual switches than traditional topologies, the fact that it only deploys commodity rack switches in the entire data center brings down the operational power costs significantly [21]. For these reasons, we expect that rack switches will be the dominant switches deployed in data center networks in the not-so-distant future.

In addition to advances in network topology design, researchers have also shown that active power management of networking devices can help reduce network energy costs significantly [9, 10, 19, 22, 29]. Such techniques can be performed both on an individual switch level as well as at the entire data center level. Individual switch power management techniques include powering off idle components such as unused linecards, and ports, as well as setting ports speeds lower during times of low traffic. Data center wide schemes include network power-aware job allocation, network power-aware routing of traffic, and simultaneously turning off unused networking elements including entire switches. Simulation results show that both energy efficient fat-tree topologies and the above power management schemes reduce operational networking power by as much as 50-70% [19, 22].

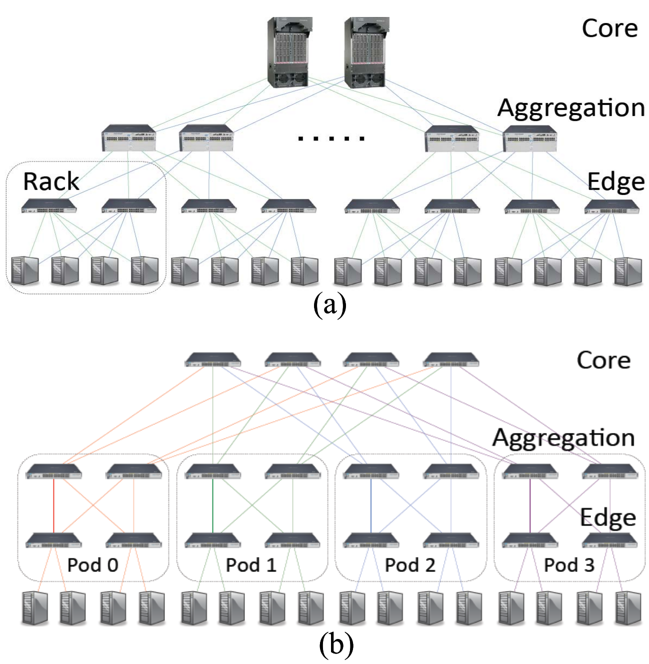

Figure 1. (a) A typical data center topology. Racks typically hold up to 40 "1U" servers and about two rack (edge) switches. (b) A fat-tree data center topology using only inexpensive commodity switches for the same number of end hosts (servers) as the typical topology.

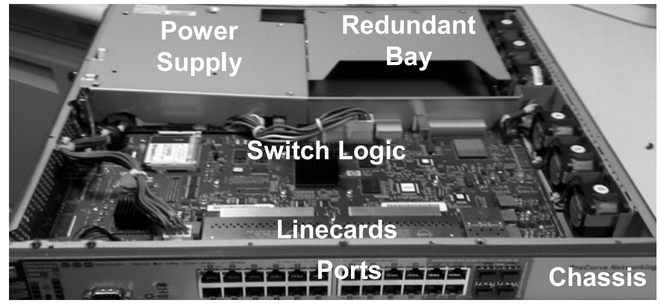

Figure 2. Typical 24-port 1U switch layout. We consider 48-port 1U switches that have similar layout in our study.

IV. Research Methods

To estimate the lifecycle energy use of a network switch, we consider three categories:

(i) energy used to manufacture the switch;

(ii) energy used in the operation of a single switch;

(iii) energy used depending on network architecture (topology), i.e. how the switches are inter-connected.

We consider three kinds of switches in this study – rack (typically 1U) switches, aggregate (typically 4U) switches, and core switches. The 1U switch that we consider in our study has a maximum switching capacity of 148 Gbps and maximum fabric speed of 173 Gbps. The switch contains 44 auto-sensing 10/100/1000 (IEEE 802.3 Type 10Base-T, IEEE 802.3u 100Base-TX, or IEEE 802.3ab Type 1000Base-T) ports and 4 additional dual-personality ports (which can be used as either RJ-45 10/100/1000 ports with Power over Ethernet (PoE) or as open mini-GBIC slots for use with mini-GBIC transceivers). The switch measured 43 cm x 44.3 cm x 4.4 cm (1U height) and can be rack-mounted. It contains two electronic modules: a management module consisting of Freescale PowerPC 8540 processor (666 MHz) with 4-MB flash, 128 MB compact flash, and 256 MB DDR SDRAM; and a 10G module containing ARM9 (200 MHz) with packet buffer size of 36 MB QDR SDRAM.

The aggregation switch sits within a 4U chassis, measuring approximately 45.09 x 44.45 x 17.53 cm, and contains 144 mini-GBIC and 10/100/1000 ports with PoE, and an additional twenty-four 10 GbE ports providing up to 214 Mbps performance. The estimated performance is 214 Mbps, with a maximum switching capacity of 322.8 Gbps on a maximum fabric speed of 345.6 Gbps. In addition to the electronic modules in the rack switch, the aggregation switch also contains a Gigabit module with an ARM9 (200 MHz) packet buffer size of 144 MB QDR SDRAM.

In addition to the above switches, we also consider the operational power use of core switches within the network. To estimate their manufacturing energy, we scale up the manufacturing energy used by the aggregation switches based on a per-unit mass estimate to obtain an approximate value of the energy used in manufacturing the core switches. While such an approach will provide only estimated values, we believe these approximations are reasonable for the purpose of the present work.

A. Energy Used During Manufacturing

We begin by disassembling an existing 1U network switch, and build a lifecycle inventory (LCI) of the different components within the switch. Table 1 summarizes the key components of the switch. For each of these components, we scale unit process data of the constituent materials and upstream manufacturing processes using the ecoinvent v2 database [23] to estimate cradle-to-gate impact. We summarize the unit data and the resulting values in Table 1.

For the aggregation switch, rather than recreate the LCI, we simply extend the LCI of the rack switch by scaling each component appropriately and adding any missing modules. For example, the aggregation switch has about three times the number of ports relative to the rack switch; therefore, we simply scale the energy used to manufacture line-cards by a factor of 3X from the rack switch to obtain the estimate of manufacturing energy in the aggregation switch. Similarly, the chassis for the aggregation switch is scaled from the rack switch estimates on a per-unit mass basis. Accordingly, we find that the energy required to manufacture the aggregation switch is approximately 67% larger than the rack switch. We similarly scale the core switch relative to the aggregation switch. These estimates should be further verified in future work.

Table 1. Life-cycle Inventory for Rack Switch.

Component	Value	Source
Power Supply	2.1 kg	Literature[a]
Printed Wiring Boards	253 in^2	Analytical[b]
Semiconductors	12.7 g	Analytical[b]
Chassis, Cables and Fans		
Steel	4.2 kg	Analytical / Literature[c]
Aluminum	5.4 g	Analytical / Literature[c]
Copper	15.1 g	Analytical / Literature[c]
Nickel	15.2 g	Analytical / Literature[c]
Zinc	10.7 g	Analytical / Literature[c]
Plastic	864 g	Analytical / Literature[c]
Transportation		
Sea	92088 kg-km	Estimated[d]
Air	72 kg-km	Estimated[d]
Road	960 kg-km	Estimated[d]

[a] Mass of the power supply is scaled from the work of Hannemann et al. [[8]].
[b] Components with 'analytical' source were measured from actual rack switch.
[c] For the chassis, the mass of the chassis was analytically measured, but the breakdown of the materials within the chassis is based on literature data published elsewhere [8][23].
[d] The transportation values were based on actual supply chain data for industry-standard servers [8] rather than the supply chain of an actual switch. However, it is likely that the two are quite similar.

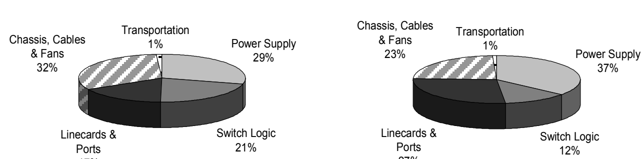

(a) (b)

Figure 3. (a) Breakdown of energy use during manufacturing of a rack switch (energy to manufacture ~ 2.8 GJ-eq). (b) Breakdown of energy use during manufacturing of an aggregation switch (energy to manufacture ~ 4.7 GJ-eq).

For the inventory shown in Table 1, impact factors from the literature [8, 23, 24] were applied to estimate the total energy used during manufacturing of the different switches. Figure 3 shows a breakdown, by component, for the rack switch and aggregation switch. We find that the energy used during manufacturing is roughly evenly split across the key components (power supply; linecards & ports; chassis; and the various logic boards). The linecards and ports, which perform the primary networking function, are responsible for less than 20% of the total rack switch manufacturing footprint. On a per-port basis, we find that the rack switch requires approximately 59 MJ-eq per port.

However, as the switch scales towards larger capacity (e.g. aggregation versus rack switch), the relative contribution from the linecards and ports becomes responsible for a larger fraction of the energy used during manufacturing. Concurrently, because the larger switches also consume more power, the size of the power supply increases; by comparison, the chassis and the switch logic boards increase by relatively lesser amounts. On a per-port basis, we find that the aggregation switch has a manufacturing efficiency of approximately 33 MJ-eq per port.

Thus, as the number of ports on a single switch is increased, the burden of manufacturing begins to shift. While more energy is spent on a per-switch basis, in terms of the fractional amount of energy being spent on 'useful' or desirable components within the switch (i.e. ports), we find that larger switches are actually more efficient from a manufacturing perspective.

B. Energy Used During Operation of An Individual Switch

Current switches do not include comprehensive operational energy consumption values. Device specification data sheets only report maximum rated power. The actual energy consumed by switches depends on various factors such as device configurations and traffic workload; thus relying only on the maximum rated power will grossly overestimate the actual energy consumption, by as much as by 70%, as reported in prior work [20].

Rather than rely upon the rated power values of the switches, we use a linear model proposed in [20] to predict the operational power of individual switches that is based on switch chassis, linecards, and number of ports as shown below:

$$P_{switch} = P_{chassis} + N_{linecards} * P_{linecard} + \Sigma_i N\_ports_i * P_i * U \quad (1)$$

where P is the instantaneous power and N is the number of instances. $P_{linecard}$ is the power consumed by the linecard with no ports turned on, and $N_{linecards}$ is the actual number of cards plugged into the switch. Variable *i* in the summation is the number of configurations for the port line rate. P_i is the power for a port running at line rate *i*, where *i* can be 10 Mbps, 100 Mbps, or 1 Gbps and *U* is a utilization scaling factor to account for the traffic through each port. The above linear model, benchmarked against actual measurements of power use in switches, is able to accurately (under 2% error) capture the power consumption of a wide variety of switches currently in use.

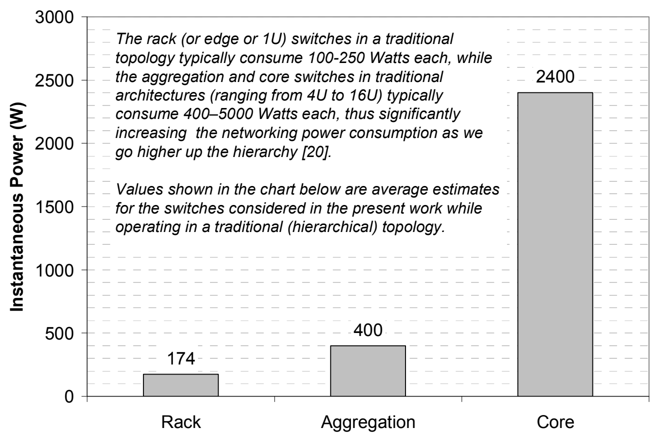

Figure 4. Instantaneous Power Use during operation of different switches.

Figure 4 shows an illustration of the instantaneous power obtained from a rack, aggregation, and core switch which is used in the present work.

C. Network Topology and Infrastructure Energy Use

The operational estimates shown in Fig. 4 above are for a typical (traditional) network topology. However, when considered holistically, the power used by an individual switch will be dependent on the topology in which the switch operates. For example, the power used by a rack switch in the traditional (hierarchical) topology of Figure (1a) will likely differ from that of the same rack switch when placed in a fat-tree topology of Figure (2) due to variations in the number of active ports/linecards, speed of the ports, traffic through the ports, etc. Furthermore, within most data center environments, the cooling infrastructure outside the switch can consume a significant amount of additional energy [8, 27]. Therefore, in addition to switch-level considerations of energy use, topology- and infrastructure-related considerations of the energy use must be included in the operating model.

Specifically, the topology and infrastructure power use can be modeled as:

$$P_{total} = PUE * (\Sigma P_{rack} + \Sigma P_{aggregation} + \Sigma P_{core}) \quad (2)$$

where P corresponds to instantaneous power. The PUE (Power Usage Effectiveness) is a measure of the infrastructure efficiency [17] and is defined in terms of the total power required within the infrastructure, per unit power provided to the IT equipment. The sum of power consumed across all the switches will be topology- and workload-dependent. Thus, Equation (2) provides an estimate of the total power consumed for networking devices within the data center. It should be noted that we do not consider the energy required to manufacture the data center infrastructure, since the lifetime of the infrastructure (20+ years) is generally much longer than that of the IT equipment (3-5 years), and a single infrastructure typically supports many thousand IT systems in each generation of equipment. Thus, we assume that on a per-switch basis, the energy required to manufacture the infrastructure is relatively small [24] and can be neglected. This assumption should be verified in future work.

D. Case Study

We consider five infrastructure scenarios, summarized in Table 2. We use traces from a subsection of a real production data center in order to calculate operational switch power accurately as well as to simulate the effects of energy management schemes to reduce operational energy. Since data center sizes vary widely, we consider two extreme scenarios – (i) a data center hosting 292 servers (since our network traces were collected from 292 servers) (ii) a data center hosting around 28,000 servers. Cases A-D involve the smaller network, supporting approximately 292 servers; while Case E involves the larger network. Case A is the traditional (hierarchical) network topology discussed earlier. In addition to this baseline case (A), we consider a scenario (B) where network administrators manage individual switches for energy efficiency by turning off switch ports and linecards when not in use. We also consider a scenario (C) where traffic is aggregated efficiently all the way from the servers to the aggregation switches. In this scenario, it becomes possible to build an efficient configuration (C1) using fewer switches; any unused switch is powered off. We keep the core switches powered on at all times, since they are critical for maintaining connectivity to the outside world Elimination of other unused switches, however, also means that the failure of any single switch in the network will disconnect the network and render servers unreachable. Therefore, one may choose instead to ensure at least one redundant path exists for traffic flow (C2). We also consider a data center network topology based on the fat-tree architecture (D) consisting of 144 1U rack switches. In such a fat-tree network, traffic is aggregated and forwarded using fewer switches, and unused switches are powered off.

For all of the above cases, using network traces from a production data center network hosting an e-Commerce application [19], we project the total lifetime energy use across all switches in the data center, assuming a 3 year lifetime and a data center PUE of 1.4. The results are shown in Figure 5. We find that for a traditional topology (A-C), the manufacturing of the network switches accounts for between 6% and 9% of the total lifecycle energy use, while the operation of the switches itself contributes the largest portion to the lifecycle energy use (between 65% and 67%) with the data center infrastructure accounting for the balance (around 26%).

Table 2. Scenarios considered in present work.

Case	*Description*	*Number of Rack / Agg / Core Switches*
A	Traditional (Hierarchical) Topology (Fig. 1a)	26/2/2
B	Traditional Topology, with unused components powered off	26/2/2
C	Traditional Topology, with Traffic Aggregation	
C1	Without Switch Redundancy	26/2/2
C2	With Switch Redundancy	26/2/2
D	Fat-Tree Topology (Fig. 1b)	72/72/36
E	Large-Scale Network	
E1	Traditional (Hierarchical) Topology (A)	2332/78/8
E2	Traditional, redundant traffic aggregation (C2)	2332/78/8
E3	Fat-Tree Topology (D)	1152/1152/576

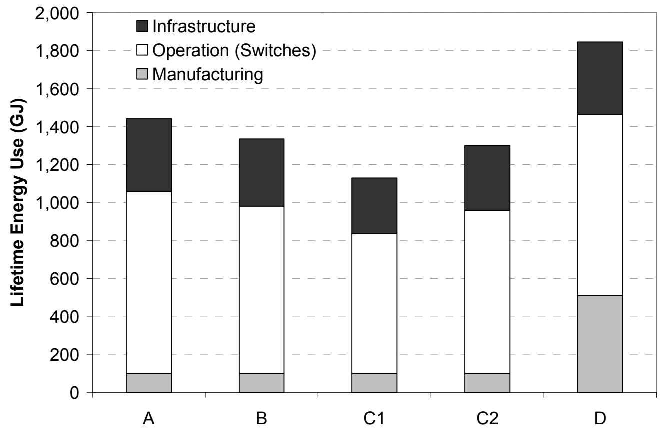

Figure 5. Lifetime energy use for some different topologies considered in Table 2, for a 292-server network.

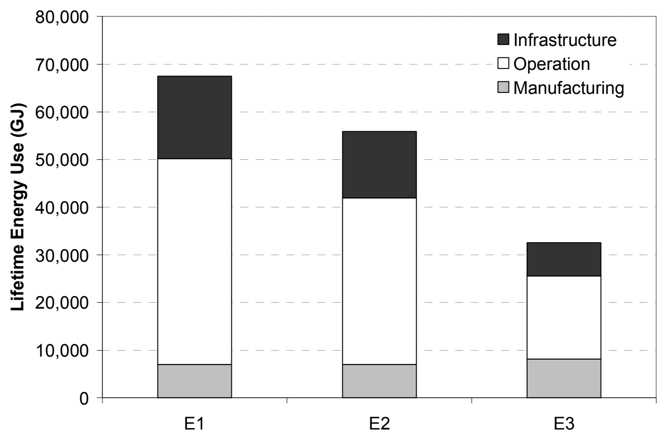

Figure 6. Lifetime energy use for some different topologies considered in Table 2, for a 1166-server network.

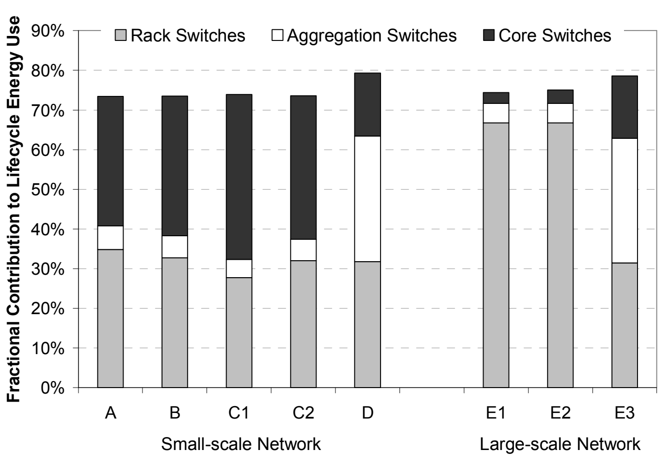

Figure 7. Relative contribution of different switch categories to the lifecycle energy use. (Balance is attributable to infrastructure energy use.) Note that for the fat-tree topologies (cases D and E3), the core and aggregation switching functions are performed by industry-standard commodity 1U switches.

For the fat tree (D), however, the large number of switches required (approximately 144 relative to around 30 for the traditional topology) result in a much larger contribution of manufacturing – nearly 28% – with the operational contribution shrinking to just below 52%. Thus, it appears that focusing around strategies to minimize the operational energy use is the most important consideration in traditional network topologies; but for emerging topologies such as the fat tree, reducing the number of switches required in the topology can help reduce the overall infrastructure burden.

Across all the cases, the total lifecycle energy use of the networking infrastructure is between 1130 GJ (for case C1) and 1845 GJ (for case D). As may be expected, relative to the base case (A), all of the power management techniques which reduce operational energy use without adding to any manufacturing overhead result in energy savings (of between 7% and 22% across the lifetime). It should be noted, however, that the most energy-efficient configuration – a traditional topology which aggregates traffic before forwarding, but without any redundant traffic path – is also the least reliable. For mission-critical enterprise workloads, case C2 (which represents about a 10% savings in lifecycle energy use relative to the base case) may be a more appropriate choice. We find that a fat-tree network (case D) is not advisable for this workload and data center size, because the network is small enough where the corresponding utilization of nodes within the traditional topology is sufficiently high that adding extra commodity switches does not provide any net benefit.

However, if a larger data center (which is also more typical) is considered (case E1-E3), we find that the benefit from a fat-tree can be significant. In this case, we scale up the traffic trace appropriately across the larger network in order to estimate the switch lifecycle energy. As shown in Figure 6, for a large-scale network, the fat-tree (E3) results in nearly 52% savings relative to the baseline topology (E1). As shown in Figure 7, most of these savings come from the ability to more efficiently aggregate and forward traffic across the different hierarchies within the fat tree, allowing unused switches to be shut off across the topology during times of low-traffic. Thus, depending on the size of the network and the traffic patterns, the optimal network topology will differ.

E. Data Quality and Uncertainty

Most of the data utilized in this case study are from the disassembly of an actual rack switch. The instantaneous power measurements of the different switches have been experimentally validated in an actual data center. Thus, most of the data utilized in the study are of fairly high quality [25]. Some uncertainty exists around the traffic patterns across the network. In addition, for the aggregation and core switches, proxy data based on the rack switches was utilized. However, a sensitivity study on feasible ranges of these data suggests that the impact in terms of overall lifecycle energy use is quite small – of the order of 1% or less.

V. Conclusions

Compared to benchmarks from previous studies [24], our analysis suggests that the networking infrastructure is responsible for between 3% and 10% of the total data center lifecycle energy use. This range agrees well with previous estimates of operational power consumed by networking equipment within data centers [27], which have estimated that between 5% and 15% of all data center electricity use could be attributed to the networking equipment.

For traditional network topologies, the largest energy use occurs during operation of the switches. Thus, methods to reduce the operational energy use of network switches such as powering off unused components, and more efficiently forwarding traffic so that unused switches can be shut off, provides a viable alternative to reducing the total energy footprint of networking within data centers.

However, the optimal choice for energy efficiency within data center networking environments will depend on the traffic patterns and size of the network. For small networks, operational energy efficiency measures such as those described above are likely to provide the greatest return. For larger networks, however, alternative network topologies such as the fat-tree network may result in the lowest energy footprint. In such fat trees, the relative share of manufacturing is much higher than traditional architectures: in other words, as network equipment becomes more operationally energy-efficient, one may expect a shifting in burden to gradually occur such that the manufacturing eventually becomes the largest contributor to the lifecycle. For such a paradigm, the key to reducing lifecycle energy use will be designing network switches with *fewer* components (rather than just more *efficient* components). One way to affect such a change may be providing increased redundancy via software management instead of hardware redundancy, as well as a more holistic approach that attempts to move certain trade-offs (such as power conversion efficiency) away from the individual switches and towards shared infrastructure. Future work will explore such considerations and trade-offs in more detail.

References

[1] S. Mingay, "IT's Role in a low carbon economy," presented at the *Gartner Symposium / IT Expo*, Cannes, France, 2008.

[2] *SMART2020: Enabling the Low Carbon Economy in the Information Age*, Report by The Climate Group on behalf of the Global eSustainability Initiative (GeSI). Available: http://www.smart2020.org/

[3] M.W. Toffel, and A. Horvath, "Environmental implications of wireless technologies: news delivery and business travel," *Env. Sci. Tech.*, vol. 38, no. 11, pp. 2961-2970, 2004.

[4] E. Williams, "Energy intensity of computer manufacturing: hybrid assessment combining process and economic input-output methods," *Env. Sci. Tech.*, vol. 38, no. 22, pp. 6166-6174, 2004.

[5] R. Kuehr, and E. Williams, *Computers and the Environment: Understanding and Managing their Impacts*. New York: Springer, 2003.

[6] Y.A. Huang, C.L. Weber, and H.S. Matthews, "Carbon footprinting upstream supply chain for electronics manufacturing and computer services," in *Proc. IEEE Int. Symp. Sust. Sys. Tech.*, Tempe, AZ, 2009.

[7] C.T. Hendrickson, L.B. Lave, and H.S. Matthews, *Environmental Life Cycle Assessment of Goods and Services: An Input-Output Approach*, Washington DC: Resources for the Future, 2006.

[8] C.R. Hannemann, V.P. Carey, A.J. Shah, and C. Patel, "Lifetime exergy consumption of an enterprise server," in *Proc. 2008 IEEE Int. Symp. El. Env.*, San Francisco, CA, 2008.

[9] K. Christensen, B. Nordman, and R. Brown, "Power Management in Networked Devices," *IEEE Computer*, pp. 91-93, August 2004.

[10] B. Nordman, "Networks, Energy and Energy Efficiency," presentation at Cisco Green Research Symposium, San Jose, CA, March 2008.

[11] J.G. Koomey, "The environmental cost of cloud computing: Assessing power use and impacts," presentation at Green:Net 09, San Francisco, CA, March 2009.

[12] D. Meisner, B.T. Gold, and T.F. Wenisch, "Powernap: Eliminating server idle power," *Proc. 14th Int. Conf. Arch. Supp. Prog. Lang. Op. Sys. (ASPLOS)*, Washington DC, March 2009.

[13] L.A. Barroso, and U. Hölzle, "The case for energy-proportional computing," *IEEE Computer*, vol. 40, no. 12, pp. 33-37, 2007.

[14] R. Bianchini, and R. Rajamony, "Power and energy management for server systems," *IEEE Computer*, vol. 37, no. 11, pp. 68-74, 2004.

[15] R. Raghavendra, P. Ranganathan, V. Talwar, Z. Wang, and X. Zhu, "No 'power' struggles: coordinated multi-level power management for the data center," *Proc. 13th Int. Conf. Arch. Supp. Prog. Lang. Op. Sys. (ASPLOS)*, Seattle, WA, March 2008.

[16] C.E. Bash, C.D. Patel, and R.K. Sharma, "Dynamic thermal management of air cooled data centers," *Proc. 10th Intersociety Conf. on Thermal and Thermomechanical Phenomena (ITHERM)*, San Diego, CA, pp. 452-460, 2006.

[17] C. Belady, A. Rawson, J. Pflueger, and T. Cader, "The Green Grid Data Center Power Efficiency Metrics: PUE and DCiE," The Green Grid, White Paper No. 6, 2008.

[18] R.K. Sharma, C.E. Bash, C.D. Patel, R.J. Friedrich, and J.S. Chase, "Balance of power: dynamic thermal management for Internet data centers," *IEEE Internet Computing*, vol. 9, no. 1, pp. 42-49, 2005.

[19] P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan, "Energy Aware Network Operations." *In Proc of IEEE Global Internet Symposium (in conjunction with IEEE Infocom)*, Rio de Janeiro, Brazil, April 2009.

[20] P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan, "A Power Benchmarking Framework for Network Devices". In *Proc of IFIP Networking*, Aachen, Germany, May 2009.

[21] M. Al-Fares, A. Loukissas, A. Vahdat., "A Scalable, Commodity Data Center Architecture". *In Proc of SIGCOMM*, Seattle, WA, August 2008.

[22] B. Heller, S. Seetharaman, P. Mahadevan, Y. Yiakoumis, P. Sharma, S. Banerjee and N. McKeown, "ElasticTree: Saving energy in Data Center Networks" *In Proc of the 7th USENIX Symposium on Network Systems Design and Implementation (NSDI)*, San Jose, CA, April 2010.

[23] *The ecoinvent v2 database*. PRé Consultants. Available http://www.pre.nl/ecoinvent/default.htm

[24] A. Shah, C. Bash, R. Sharma, T. Christian, B.J. Watson, C. Patel, "The environmental impact of data centers," *Proc. IPACK 2009*, San Francisco, CA, July 19-23, 2009.

[25] B.P. Weidema, Mand .S. Wesnaelig, "Data Quality Management for Life Cycle Inventories – An Example of Using Data Quality Indicators," *J. Cleaner Prod.*, vol. 4, no. 3, pp. 167-174, 1996.

[26] A. Greenberg, J. Hamilton, D.A. Maltz, P. Patel, "The Cost of a Cloud: Research Problems in Data Center Networks." CCR January 2009, Editorial Note.

[27] "Report to Congress on Server and Data Center Energy Efficiency:" US Environmental Protection Agency, EnergyStar, http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf, August 2007.

[28] J. J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, "Power Awareness in Network Design and Routing." In *IEEE-INFOCOM*, April 2008.

[29] S. Nedevschi, J. Chandrashenkar, B. Nordman, S. Ratnasamy, and N. Taft. "Skilled in the Art of Being Idle: Reducing Energy Waste in Networked Systems." In Proceedings Of NSDI, April 2009.