Energy Proportionality of an Enterprise Network

Priya Mahadevan, Sujata Banerjee, and Puneet Sharma
HP Labs
Palo Alto, California, U.S.A
priya.mahadevan@hp.com, sujata.banerjee@hp.com,puneet.sharma@hp.com

ABSTRACT

Energy efficiency is becoming increasingly important in the operation of networking infrastructure, especially in enterprise and data center networks. While strategies for lowering the energy consumption of network devices have been proposed, what is lacking is a comprehensive measurement study conducted across a large network (such as an enterprise), that monitors power usage as a function of traffic flowing through the network. We present a large power profile study that we conducted in an enterprise network, comprising of 90 live switches from various vendors. We first describe Urja, the system that we built, that collects required configurations from a wide variety of deployed switches and uses them to accurately predict the power consumed by individual devices and the network as a whole. Urja is vendor neutral, and relies on standard SNMP MIBs to gather the required configuration and traffic information. Further, based on available knobs in current devices, the analysis engine in Urja lists various configuration and rewiring changes that can be made to the devices in order to make the network more energy proportional. Urja has been deployed in an enterprise sub-network for about 4 months; through comprehensive analysis of the data collected over this period, we present various changes (in increasing order of cost and complexity) that network administrators can perform; in this segment of an enterprise network, we can save over 30% of the network energy through simple configuration and rewiring changes, and without any performance impact.

Categories and Subject Descriptors

C.2.3 [**Network Operations**]: Network management; C.2.3 [**Network Operations**]: Network monitoring

General Terms

Measurement, Management

Green Networking 2010, August 30, 2010, New Delhi, India.

Keywords

Network power, enterprise networks

1. INTRODUCTION

Energy efficient infrastructures or green IT has recently become a hot button issue for most corporations as they strive to eliminate every inefficiency from their enterprise IT systems and save capital and operational costs. Vendors of IT equipment now compete on the power[1] efficiency of their devices, and as a result, many of the new equipment models are indeed more energy efficient. However, compared to other IT devices such as servers and laptops, energy efficiency of networking equipment has only recently received attention since networks, being a shared resource, are expected to be always on. However, power consumed by the network is significant and growing. Various studies have estimated the annual electricity consumed by networking devices in the U.S. in the range of 6 - 20 Terra Watt hours [13, 14].

In addition, the following three trends are motivating researchers to address power management in networks:

- The rich work done on server power management and cooling technology over the past several years is causing the network power consumption to be a bigger fraction of the overall power budget. The next natural frontier of IT power savings is in the network area.
- New technologies that demand higher network bisection bandwidth require large power hungry network switches. From one study [1], the network switch to connect 1000 10GbE hosts with a 1:3 bisection bandwidth between the nodes would require about 40 watts per edge NIC.
- IT workloads vary over time and are increasingly being consolidated over virtualized infrastructures to a minimal set of physical servers leaving many network devices to be idle.

Researchers have recently proposed various network-wide energy management schemes for deployment in a large data center or wide area network. However, the huge legacy base of IT equipment that will be in the system for some time to come also needs attention with respect to the energy efficiency issues. In this paper, we build on our prior work on network power benchmarking [9], network power savings

9[1]We use energy and power interchangeably in this paper.

in data center networks [8, 10] and apply some of the ideas to enterprise networks. However, enterprise networks are inherently different from data center networks.

One of the biggest hurdles in making enterprise networks more energy efficient is that these networks have a higher diversity of devices from multiple vendors with respect to both the models and the age of the devices. Many enterprise networks grow organically with ultimate topologies that are hard to operate efficiently with respect to energy. Shrinking IT budgets may cause many of these devices to operate for as long as possible without the possibility to replace them with newer lower power models. Enterprise network management is still not fully automated and typically requires per device manual configurations that may be error prone and unlikely to make informed decisions for network-wide energy efficiency. One of our goals is to operate networks power proportionally, i.e., consume power in proportion to the load. The challenge is that individual network devices today are far from being power proportional [9].

In this paper, we describe a network wide energy monitoring tool that we built called Urja. Urja collects configuration and traffic information from live network switches and accurately predicts their power consumption. By analyzing real network traces, we provide several techniques that can be integrated into network management operations so as to get significantly closer to power proportional behavior with today's non power proportional devices. Further, we discuss the practical aspects of implementing these techniques in any enterprise network. To our knowledge, ours is the first such large scale power study of an enterprise network.

In the next section, we provide background on network power modeling. We describe Urja, the scalable network monitoring framework, that can be deployed in an enterprise network in Section 3. Section 4 contains the characteristics and analysis of an enterprise (sub)network obtained from measurements of 90 switches. In Section 4.1, we describe a series of network management techniques that can be applied to reduce the energy footprint and present results on how much savings in energy can be obtained from each step. Section 6 summarizes our findings and presents future directions.

2. BACKGROUND

In [9], we conducted a detailed power benchmarking study of a variety of network devices ranging from wireless access points to edge LAN switches to high-end switches/routers. We studied the power consumption of individual devices both as a function of traffic flowing through them as well as their configurations. Further, we explored both the energy efficiency of switches (joules expended to transmit 1 Mbps of traffic) as well as their energy proportionality (whether amount of energy consumed is proportional to the traffic forwarded).

We proposed a model to predict the total power consumed by the switch [9]; we find that a linear model is able to accurately (within 2% error margins) capture the total power consumption of switches/routers currently in use. As new architectural and design changes are implemented in these devices, a linear model might not be the best fit; we might have to develop other models in the future. The power consumed by a switch is given by

$$P_{switch} = P_{chassis} + num_{linecards} * P_{linecard} + \sum_{i=0}^{configs} numports_{configs_i} * P_{configs_i} * UF_i$$

$P_{linecard}$ is the power consumed by the linecard with all ports disabled, and $num_{linecards}$ is the number of active cards in the switch. Variable *configs* in the summation is the number of configurations for an enabled port. $P_{configs_i}$ is the power for a port operating at speed i, where i can be unused, 10 Mbps, 100 Mbps, 1 Gbps, etc. and UF_i is the scaling factor to account for a port's utilization (traffic through a port). For details, please refer to [9].

Based on its administrative status, a port can either be enabled or disabled. A disabled port does not consume any power. Any port that is enabled consumes power. Operationally, an enabled port can either be unused i.e. have no client/cable connected to it or it can be active and capable of forwarding traffic at its set maximum capacity (10 Mbps, 100 Mbps, 1 Gbps, etc.). The energy consumed by an enabled port depends on its operational status. From our study, we summarize the various knobs that can be tuned to reduce the energy consumption of these switches and to make the networks more energy proportional: (i) turn off (disabling) unused ports, which can save up to 0.5 W per port. (ii) rate adapt a port to a lower speed based on traffic flowing through it; possible savings of 0.4 - 1 W per port when a 1 Gbps port is set to either 10 Mbps or 100 Mbps. (iii) if supported by the switch, power off a linecard if all ports on the card are disabled for savings of 30 - 60 W per card. (iv) power off an unused switch, where depending on the switch model, this action will save about 100 - 2000 W. Based on the switch power prediction model, we have built a large-scale power monitoring and management system for enterprise networks. We describe this system in the next section.

3. POWER MONITORING FRAMEWORK

We show the architecture of Urja, our power monitoring and management tool in Figure 1. Urja has 4 separate components - the Measurement based switch power model, Web-based power profiler, Analysis engine and Power management engine. Urja has a database that stores the power constants associated with all switch models, line card types, etc. The web-based power profiler polls all (or a subset) of switches in a network and obtains relevant configuration information from them using standard entity MIBs over SNMP. The information that is polled from the switches include the switch chassis type, firmware version, number and type of active linecards, number of active ports on each card, administrative status of each port (enabled or disabled), operational status of each port and the traffic flowing through each port. Based on this information, the web-based profiler, uses the appropriate power constant values from its database and uses the switch power model to predict the power consumption of the switch. Urja displays the real-time power consumption of each switch in the enterprise network, along with the total power consumed across all the switches. The Analysis engine analyzes the data gathered from each switch and correlates this configuration information to the power consumed by the switch and the traffic flowing through it. It then lists various suggestions that can be implemented network-wide by the administrators in order to save energy and make the network more energy proportional. The Power management engine can be used by the administrators to incorporate some of the suggested configuration changes on the switches.

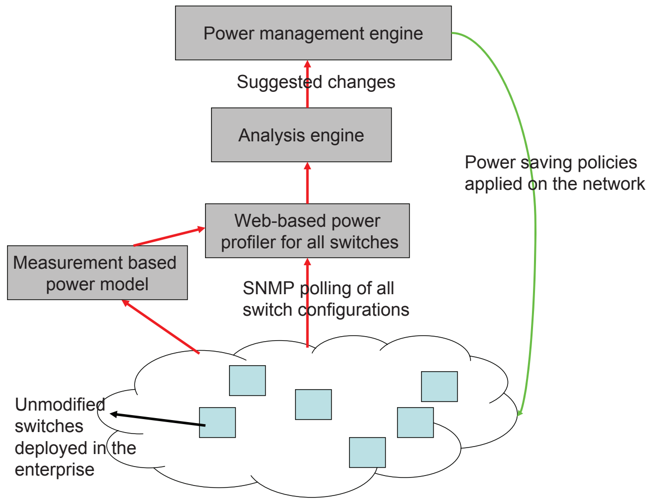

Figure 1: Urja architecture

4. ANALYSIS AND RESULTS

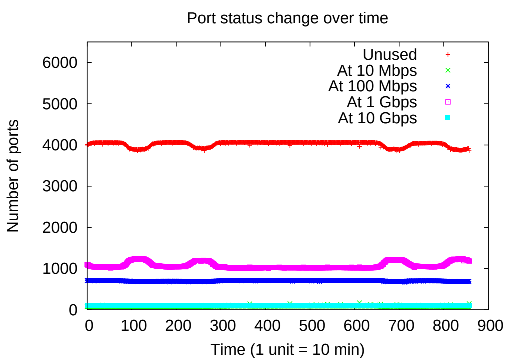

Figure 2: Variations in operating speed of ports observed in our traces as a function of time.

Urja has been deployed over a segment of an enterprise network and has been monitoring 90 switches for over 4 months. Since we have analyzed the power profiles for only a few switch models at this time, we restricted the set of switches monitored to only those for which we had power profiles. Thus the set of switches chosen may not cover complete subnets or complete sub-topologies and we plan to extend our study to larger number of switch models in the future. Of these 90 switches, about half connect to employee offices and conference rooms while the other half connect to servers in racks in server rooms. Using the results from our tool, we have identified several network operations and management techniques to make the enterprise network more energy proportional. While different enterprise networks may use different operational policies, we have reasons to believe these findings will be useful for making other enterprise networks more energy proportional as well.

In Table 1, we provide an overview of the switches in the chosen enterprise (sub)network. We list switch chassis power, and power of an individual port in the various operational states. For switch model B that supports pluggable

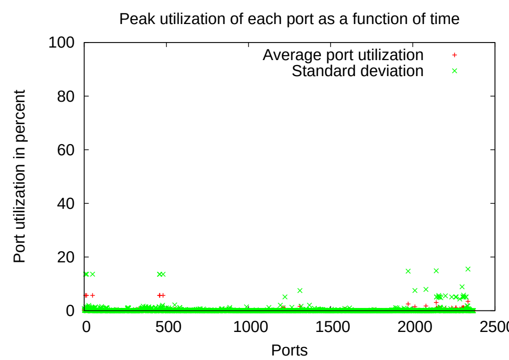

Figure 3: Percent utilization ((averaged over 6 days) for all active ports.

linecards, there is an additional fixed power cost for each linecard plugged in (30 - 38 W depending on the card type). The number of ports on each switch depends on the linecards plugged in and the number of ports on each card. Power cost of the switch chassis and individual ports vary depending on other factors such as firmware version, linecard type etc. We do not provide an exhaustive list of power values for each case in Table 1; the power model in Urja stores all these details and uses the appropriate values to compute the network power accurately. The power values for all active ports in 10/100/1000 Mbps mode correspond to 0% port utilization; a significant increase in port utilization increases the power cost per port slightly. The 90 switches are diverse in terms of size (number of linecards, ports etc.) and are from different manufacturers and product generations. We believe this is quite typical in many enterprise networks. The total number of ports across these 90 switches is 6710 - 6 switches have 32 ports each, 19 switches have 48 ports each, 3 switches have 50 ports each, 19 switches have 52 ports each, 5 switches have 76 ports, 26 switches have 100 ports each and 12 switches have 124 ports each.

In March 2010, over a 6 day period, we recorded switch

Switch Model	Number of Switches	Supports pluggable linecards?	Chassis power in Watts	Power per enabled but unused port in Watts	Power per 10 Mbps port in Watts	Power per 100 Mbps port in Watts	Power per 1 Gbps port in Watts
A	20	No	70	0	0.2	0.3	0.65
B	14	No	147	0	0.12	0.18	0.89
C	53	Yes	55	0.25	0.4	0.48	0.9
D	3	No	50.3	0.5	0.5	0.5	0.5

Table 1: Overview of switch type and some of their power constants for the enterprise network.

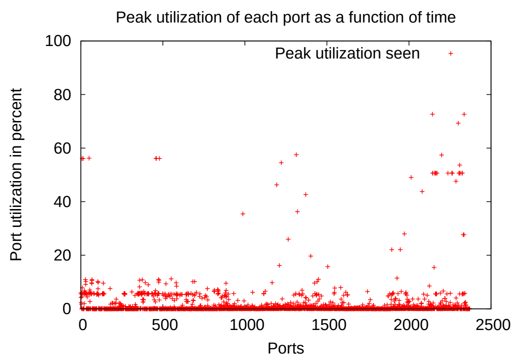

Figure 4: Peak utilization in percent observed for all active ports.

configuration and traffic through these 90 switches in a fine grained timescale and derived the estimated power consumed using Urja. Our trace starts on Wednesday, March 3, 6:30 pm and ends on Tuesday, March 9, 5:00 pm and switches are polled every 10 minutes. Since our trace period includes both weekdays and weekends, we see differences in port operational status and the corresponding power variations during this period.

Of the 6710 ports, 716 have been explicitly disabled by the network administrators and thus do not consume any additional power. The remaining 5994 ports are enabled and consume power, irrespective of whether they are used to forward traffic or not. While the operational status of a port varies during our trace duration, we report port operation status at the start of our trace duration. 86 ports were operating at 10 Mbps capacity, 707 ports at 100 Mbps, 1150 ports at 1 Gbps capacity and 106 ports at 10 Gbps. The number of ports that are enabled but unused at the start of our trace is 3945.

During our trace period, no administrative change was performed on the port status, i.e. no disabled port was enabled by the administrator and vice versa. Only changes that occurred were in the operational status of the ports. In Figure 2, we plot the changes in the operational status of the (enabled) ports that we observed in our traces as a function of time. The number of ports at 10 Mbps, 100 Mbps and 10 Gbps remains relatively constant during the entire trace period. The number of ports at 10 Mbps and 10 Gbps are almost the same, and the lines in the plot overlap. The most variations are observed for ports operating at 1 Gbps and ports that are unused. An increase in the number of ports at 1 Gbps also corresponds to a dip in the number of unused ports. In our trace, bumps in the 1 Gbps line corresponds to day time - usually between the hours of 8 am and 6pm, while the relatively flat portion of the same line corresponds to nights and weekend. On finer examination of the traces, we find that these ports typically connect to employee offices or conference rooms. We clarified that administrators did not change the speed of any of the ports during the trace duration. All port status changes occurred due to devices being plugged into the ports and the auto-negotiated rate between the device and the port. When a laptop is plugged into an office or conference room port, an unused port becomes active and operates at the auto-negotiated rate (typically 1 Gbps).

We observe that 3134 enabled ports were unused during our entire trace duration. Apart from these unused ports, we found 488 ports had zero utilization throughout the trace duration. These 488 ports were enabled and connected to devices, though we observed no traffic on these ports during the 6 day period. It is likely that these ports are connected to servers or desktops that were shutdown and never used during those 6 days. Analyzing all the traffic through the individual ports across the trace duration (Figure 3), we find that all active ports have an average utilization of under 10%. The standard deviation of the average utilization for each of these ports is under 20. In Figure 4, we plot the peak utilization that we observed in our 6 day period for each port. Though anecdotal, we believe that this data is representative of many other enterprise network environments.

Maximum observed total power consumed across these switches is 18370 Watts while the minimum total network power is 18190 Watts. Average network power consumption during the 6 day period is 18229 Watts. The small variations that we see in power are due to the variations in port operational status such as an enabled but unused port getting active at, say, 1 Gbps. In Figure 5, we plot the the total traffic through the network on the y-axis and the total network power consumed on the secondary y-axis. We note that the variations in power are minor and is influenced by the number of ports that undergo operational status change. The average utilization of each port is low and has negligible impact on the power consumption. The peaks in the traffic lines correspond to backup traffic that happens at night time, while the peaks in the power curve occur when the office and conference room ports become used (active) again. In fact, the shape of the power curve in Figure 5 follows the shape of the 1 Gbps line in Figure 2.

4.1 Possible techniques to make the network more energy proportional

Given the above trends in the network utilization, our goal is to exploit them and incorporate practices to start saving energy in the existing legacy base in enterprise networks.

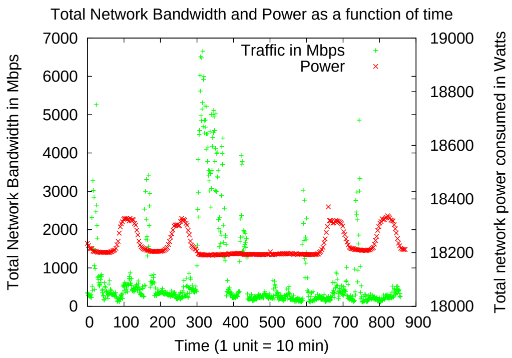

Figure 5: Total traffic as a function of time. On the secondary axis, we plot total network power.

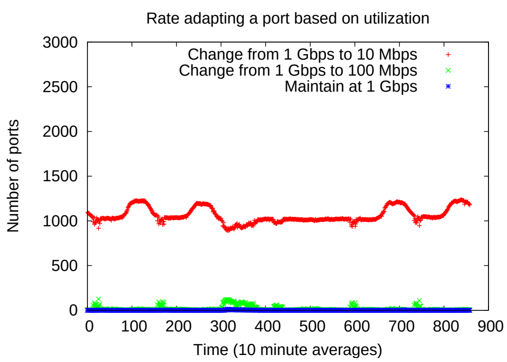

Figure 6: Port rate adaptation

While we present results for a 6 day period, longer duration such as monthly analysis of the same metrics are interesting, which is a topic of future study. In this section, we present a sequence of steps that can be employed by network administrators to make their infrastructure more energy efficient.

Technique 1: Disabling unused ports

Powering off (disabling) unused network ports is an easy first step to saving energy. The cost of implementing this step is extremely low, as network administrators only need to turn off the ports that are unused. Of all the enabled ports in our network that we analyzed, we find that 3134 ports were never used (no cables were connected to these ports) during the entire trace duration. Further, the number of partially used ports varied during the trace duration; on an average, 4018 of the ports are not used at any given time. As we noted before, many ports in our trace are used for only short periods of time. These ports have been enabled by the administrator, and continue to draw power even when they are unused (idle). Of the 3134 never used ports, 3065 belong to switch model C and 9 belong to model D (Table 1); by disabling these ports, we can save 770 W in the enterprise (sub)network. Over a whole year, this saving translates to 6745 kWh (kilowatt-hour) of electricity.

The disadvantage of this technique is that administrators now need to enable the port using SNMP or command line interface from network management tools, before using a disabled port. Further, we observe that a relative large number of ports are only used during the day time (8am to 6 pm). Though they have been enabled, hosts are connected to these ports only during the day. Fine grained analysis of the trace show that these ports connect to employee offices and conference rooms. Network administrators can set either time-based policies, where these ports are enabled at 7 am everyday and disabled at 7 pm everyday, or use other mechanisms such as in-room occupancy sensors, or conference room booking systems to enable or disable these ports.

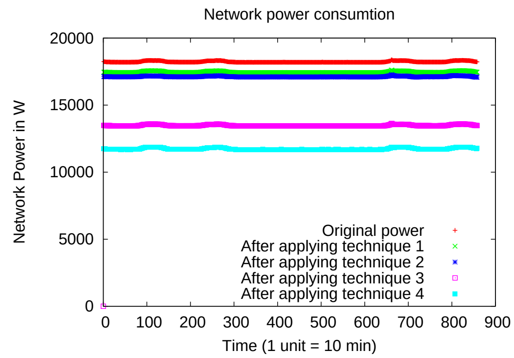

Figure 7: Network power consumption after applying the 4 energy savings techniques.

Technique 2: Port rate adaptation

Based on port utilization numbers, we find individual ports that can be set at a lower speed. Ports that have very low average utilization and don't exhibit much variability in utilization over time are good candidates to be set at a lower speed. This technique is only feasible for ports operating at 1 Gbps (they can be set to either 10 Mbps or 100 Mbps); with current technology, it is not feasible to rate adapt 10 Gbps ports. Setting a port to a lower speed may increase latency and queuing delays for packets forwarded by that port. One way to mitigate these effects is to add a slack capacity while adapting a port's speed. Thus instead of adapting a port's speed to the lowest value permitted by the traffic, we ensure that a port's utilization at that time is always less than the slack factor times its new speed, else we set the port speed to the next higher value permitted. Slack factor is configurable and can be determined by the administrator. Figure 6 shows the new speeds at which all 1 Gbps ports can be set, as a function of varying traffic through the 6 day period when we use 60% as the slack factor i.e., we ensure that the port's utilization is always less than 60% of its new speed. Very few ports need to be operated at 1 Gbps as shown in Figure 6; most ports can be set at 10 Mbps or 100 Mbps. The time taken to rate adapt a port is usually of the order of 1-2 seconds depending on the device. Thus, port rate adaptation is only suitable at a coarser time granularity leveraging long term trends over hours or even days. Further, it is unreasonable to expect administrators to manually adjust port speeds throughout the day. Instead, this process needs to be automated. Given the fact that power required to run a port at 100 Mbps instead of 10 Mbps is extremely low, all ports in Figure 6 whose new speed varies between 10 Mbps and 100 Mbps over time can instead all be set at 100 Mbps. By taking this action, the average network energy over the 6 day period is 17413 W, for an average savings of 816 W. This technique can be easily combined with

Technique 1; combining Techniques 1 and 2 in the enterprise (sub)network saves us, on an average, 1586 W equivalent to 9% energy savings.

IEEE's 802.3az standards aims to bring more energy proportionality in a port's energy consumption. We believe there will be no explicit need to perform this task for future energy efficient devices that will support 802.3az. However, this technique will still be useful for legacy devices. Cost of this technique is relatively low - network administrators can write scripts or policies that reduce a port's speed based on its historic utilization. This technique is not advisable for ports whose utilization shows a lot of variability.

Technique 3: Maximizing active ports on a linecard Based on the network configuration information in our traces, we observe that in most switches, a single linecard is sufficient to support all the active ports in that switch. We observe that 125 linecards (each consuming 38 W) across 48 switches of model C can be completely disabled in the observed enterprise (sub)network of 90 switches. Maximizing the number of active ports on a linecard and using the minimal possible linecards on a single switch yields a savings of $125 * 38 = 4750$ W for the enterprise (sub)network, which translates to about 26% energy savings. We note that these 48 switches support disabling unused linecards. Other switches do not support this feature, else the savings can be even greater. This technique requires rewiring on each switch, and is thus more expensive but easy to instill as a best practice.

Technique 4: Using fewer switches
Same number of active ports in the network can be consolidated across fewer switches through smarter use of the switches, though this requires significant rewiring, which is expensive. By rewiring ports from a sparsely populated switch on to a switch that has a few spare ports, one can completely power off the sparsely populated switch. Such consolidation, however, needs careful thought and planning. We need to ensure that we do not negatively impact the intra-switch capacity, inter-switch capacity, as well as the reachability (connectivity) across all the ports in the network.

To implement this technique, we consider the switch with least number of active ports and consolidate its ports across another switch (if possible) that can accommodate these ports. We ensure that we respect the different types of linecards that both switches have and do not consolidate them if there is a mismatch in card type. We also only consolidate one switch on to another switch, as without this check, we might not have the same reachability across all ports in the consolidated switch. By choosing the switch with the next lowest number of active ports, we repeat the above steps; we find that we can reduce the number of active model C switches to 27. Such a scheme is possible in the enterprise (sub)network since switch model C has a backplane capacity that can handle the extra ports consolidated on to it. However, without information about the network topology, intra-switch bandwidth cannot be guaranteed as in the pre-consolidation phase. Since in this segment of the enterprise network, average utilization of individual ports is very low, we ignore the intra-switch capacity. Using this technique, we are able to save 6233 W of network power or 34% energy savings. We note that this consolidation is not the optimal; if we had the actual enterprise topology, we could formulate it as an optimization problem. Determining the most energy efficient topology using fewest switches is for future work. In Figure 7, we plot the actual network power consumption as well as power consumption after each of our 4 techniques are applied. Technique 4 saves the most energy but is also the most complex; while Technique 1 saves the least energy, but is easiest to incorporate.

5. RELATED WORK

One of the earliest to propose energy management for networking were Gupta *et al.* [6]. Since then, researchers have proposed techniques such as putting idle sub-components to sleep [5, 6, 7, 12], as well as adapting the rate at which switches forward packets depending on the traffic [4, 5, 6, 7] and discussed their feasibility. Nedevschi *et al.* [12] propose shaping the traffic into small bursts at edge routers in order to allow network devices to sleep and rate adapt, and thus save energy. A more recent work [11] discusses the benefits and deployment models of a network proxy that would allow end-hosts to sleep while the proxy keeps the network connection alive. Chabarek [2] *et al.* enumerate the power demands of two routers; further the authors use mixed integer optimization techniques to determine the optimal configuration at each router in their sample network for a given traffic matrix. Researchers have also proposed energy management of networking devices in environments such as data centers [3, 8, 10]. However, what has been lacking so far is a large-scale power measurement study taken from a real enterprise network and correlating the power consumption with traffic flowing through various ports. Based on such a study, we propose certain best practices as well as configuration and rewiring changes that network administrators can act on in order to make an enterprise network more energy efficient and energy proportional.

6. IMPLICATIONS AND CONCLUSIONS

We present a large scale power study of an enterprise (sub) network spanning 90 switches. We record the configuration data of the switches, monitor the traffic through each port and estimate the power consumed by all the network devices using previously developed network device power models. We then present several steps that network managers can take today to start saving energy. Our goal is to exploit traffic trends and use device specific power saving knobs such as disabling unused ports, rate adapting ports, turning off entire line cards and switches as appropriate to save energy. We also present results on the potentially higher energy efficient but more expensive technique of rewiring the network. To summarize our results, the overall energy consumption can be reduced by up to 36% of what was being consumed before our network power management steps. In an enterprise setting, we find that a large number of devices are under-utilized and as expected the traffic patterns follow the patterns of an employee work day. We are exploring methods to detect usage patterns (for example, conference room bookings, room occupancy sensors, employee entry and exit, etc.) to drive just-in-time turning on of appropriate network ports and indeed entire devices. Our ultimate goal is to make networks power proportional even before all devices over time become power proportional.

7. REFERENCES

[1] http://www.voltaire.com/assets/files/SOL-BRIEF-10GB-eth-8500-WEB.pdf.

[2] J. Chabarek, J. Sommers, P. Barford, cristian Estan, D. Tsiang, and S. Wright. Power awareness in network design and routing. In *Proceedings Of INFOCOM*, April 2008.

[3] G. Chen, W. He, J. Liu, S. Nath, L. Rigas, L. Xiao, and F. Zhao. Energy-aware server provisioning and load dispatching for connection-intensive internet services. In *NSDI*, April 2008.

[4] C. Gunaratne, K. Christensen, B. Nordman, and S. Suen. Reducing the Energy Consumption of Ethernet with Adaptive Link Rate (ALR). *IEEE Transactions on Computers*, 57:448–461, April 2008.

[5] M. Gupta, S. Grover, and S. Singh. A feasibility study for power management in LAN switches. In *IEEE ICNP*, October 2004.

[6] M. Gupta and S. Singh. Greening of the Internet. In *Proceedings Of ACM SIGCOMM*, August 2003.

[7] M. Gupta and S. Singh. Energy conservation with low power modes in Ethernet LAN environments. In *Proceedings Of IEEE INFOCOM (MiniSymposium)*, May 2007.

[8] B. Heller, S. Seetharaman, P. Mahadevan, Y. Yiakoumis, P. Sharma, S. Banerjee, and N. McKeown. ElasticTree: Saving Energy in data Center Networks. In *NSDI*, April 2010.

[9] P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan. A Power Benchmarking Framework for Network Devices. In *Proceedings of IFIP Networking*, May 2009.

[10] P. Mahadevan, P. Sharma, S. Banerjee, and P. Ranganathan. Energy Aware Network Operations. In *Global Internet Symposium*, May 2009.

[11] 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 *NSDI*, April 2009.

[12] S. Nedevschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall. Reducing Network Energy Consumption via Rate-Adaptation and Sleeping. In *Proceedings Of NSDI*, April 2008.

[13] B. Nordman. Networks, energy and energy efficiency. Cisco Green Research Symposium, March 2008.

[14] K. Roth, F. Goldstein, and J. Kleinman. Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings Volume 1: Energy Consumption Baseline. *National Technical Information Service (NTIS), U.S. Dept. of Commerce, Springfield, VA 22161*, 18(NTIS number: PB2002-101438), 2002.