

A Power Benchmarking Framework for Network Devices

Priya Mahadevan, Puneet Sharma, Sujata Banerjee,
and Parthasarathy Ranganathan

HP Labs

1501 Page Mill Road, Palo Alto, CA 94306, USA

{priya.mahadevan,puneet.sharma,sujata.banerjee,partha.ranganathan}@hp.com

Abstract. Energy efficiency is becoming increasingly important in the operation of networking infrastructure, especially in enterprise and data center networks. Researchers have proposed several strategies for energy management of networking devices. However, we need a comprehensive characterization of power consumption by a variety of switches and routers to accurately quantify the savings from the various power savings schemes. In this paper, we first describe the hurdles in network power instrumentation and present a power measurement study of a variety of networking gear such as hubs, edge switches, core switches, routers and wireless access points in both stand-alone mode and a production data center. We build and describe a benchmarking suite that will allow users to *measure and compare* the power consumed for a large set of common configurations at any switch or router of their choice. We also propose a network energy proportionality index, which is an easily measurable metric, to compare power consumption behaviors of multiple devices.

Keywords: Network energy management, Benchmarking.

1 Introduction

Energy efficiency has become crucial for all industries, including the information technology (IT) industry, as there is a strong motivation to lower capital and recurring costs. According to recent literature, the annual electricity consumed by networking devices in the U.S. is 6.06 Terra Watt hours, which translates to around 1 billion US dollars per year [1], thereby presenting a strong case for reducing the energy consumed by networking devices such as hubs, access points, switches and routers.

Unlike wireless networks, energy management¹ for networking devices such as hubs, switches and routers in wired networks has not received much attention until very recently. Researchers have proposed several strategies to make routers

¹ We use power and energy management interchangeably in this paper. There is a distinction between power management for heat density versus electricity costs; however in this paper, we do not distinguish between these two issues.

and switches more energy-aware such as link rate adaptation during periods of low traffic and sleeping during no traffic [1–4]. Currently, however, there are no hardware implementations of the same. More importantly, the exact energy savings by adopting these techniques are poorly understood. One impediment to innovation in this area is the lack of power measurements from live networks and a good understanding of how the energy consumed varies under different traffic loads and switch/router configuration settings. There are many vendors manufacturing a wide variety of network devices and as of yet, there has been no focus on standardized benchmarks to measure the energy consumption of these devices. We also need benchmarks to compare the effectiveness of various energy efficient improvements.

In this paper, we attempt to fill this gap by quantifying the energy consumed by a wide variety of networking gear ranging from core switches to wireless Access Points (APs) from three different networking vendors, and under different traffic and network configurations. Our goal is to build a publicly available benchmarking suite that can be deployed on any switch; our suite sets various configurations in the switch under test and generates different traffic patterns, while a power-meter connected to the switch measures the average power for the duration of the experiment.

Non-uniformity across switches from different manufacturers, and different functionality available in the switches, make the task of identifying the right set of configurations to include in the benchmark suite challenging. Additionally, we would like our benchmarking suite to be modular, so that new components for power instrumentation can be easily incorporated. One additional challenge for device manufacturers is to ensure that networking devices such as hubs, switches and routers are energy proportional, *i.e.* they consume energy proportional to their load, similar to energy proportional laptops and servers [5, 6]. In this paper, we also discuss how energy proportionality can be measured for networking devices and their components. From our initial benchmarking study, we find that:

- There is great variability amongst switches with respect to their maximum rated power. The ratio of the actual power consumed by the device on an average, to its maximum rated power varies widely across different device families. Thus, relying merely on the maximum rated power can grossly overestimate the total energy consumed by these networking equipment.
- Energy consumed by a switch increases linearly with the number of linecards plugged into the switch as well as the number of active ports on each card.
- Energy consumed by a switch is largely independent of the packet size for a fixed traffic throughput.
- Ideally, devices should consume energy proportional to their load [5, 6]. To quantify this behavior, we define an energy proportionality index (EPI) for network devices (Section 3). The EPI values for the evaluated devices clearly demonstrate non-energy proportional behavior.

The rest of this paper is organized as follows: We begin with describing challenges in obtaining large-scale power measurements in Section 2. Our methodology to characterize power consumption in switches and routers is described

in Section 3. In Section 4, we present results from our suite. We discuss related work in Section 5. We conclude with the implications of the results of our benchmarking study in Section 6.

2 Challenges

Current switches and routers don't include comprehensive energy consumption values. Device specification data sheets only report maximum rated power. This value by itself is insufficient to understand the actual energy consumption of the networking device. As we show in the rest of this paper, the actual energy consumed by switches and routers depends on various factors such as device configurations and traffic workload; thus relying only the maximum rated power will grossly overestimate the actual energy consumption, by as much as by 70%, in some of our observations.

Power instrumentation of a device, in a live network deployment, requires unplugging the device from its power supply and plugging it into a power meter. Since the device will not be functional for several minutes, this task needs to coincide with regular network downtimes. Along with the power measurements, one also needs information on exact device configurations as energy varies significantly with different configuration settings. One also needs to record actual traffic traversing through the device and measure the energy consumed at regular intervals.

Most existing networks are not equipped with internal power instrumentation to provide information on how much energy is being consumed by each component at any given time. Adding external power instrumentation for each network device in operational networks is quite cumbersome. This task involves working with network operators to find the appropriate opportunities such as maintenance downtime to install power meters in-line with the network equipment power cables. Eventually, we expect more support for internal power measurements in the next generation of network equipment (much like the battery meters on laptops). Ideally, such power measurements need to be incorporated into standardized SNMP MIBs so that they can be queried by management applications. However, the large base of legacy network equipment deployed today does not have such instrumentation built in. Motivated by these facts, we present a model-based power inference mechanism that can be scaled up to large networks.

The idea is to generate power consumption models of networking equipment using standard benchmarking. Such models have already been defined for servers (for instance see [7]). We describe such a benchmarking suite in the next section. Using the measurements obtained from our benchmarking suite, we build a model to predict the power consumed by any networking switch or router; the device's configuration and traffic flowing through it are specified as input to the model. Once the models are available, using standard SNMP MIBs or command line interface (CLI) query mechanisms, the network operator can query the configuration and traffic information from the switches/routers and plug these values

into our models to derive the power consumed by a single box, multiple boxes and the whole network. These models can be also used in simulation/emulation testbeds. While models can perhaps not take the place of actual in-line power measurement, based on our small study, we believe that it is fairly effective, and certainly much better than using the rated (plate) power ratings provided by network vendors.

3 Benchmarking Framework

We begin by describing factors that are commonly used and most likely to affect a switch or router's power consumption. Next, we describe our efforts in building the benchmark suite, along with results from a few selected switches/routers that span from low end devices to ones that are deployed in high traffic data centers and network backbones.

3.1 Benchmarking Suite

Currently, there are no standard benchmarks used in power measurement studies of network components. One of our contributions is to develop such a benchmark that can be used to compare the power characteristics of a variety of network devices. Figure 1 shows the architecture of our benchmarking suite. The network device to be benchmarked is connected to the power outlet via a *power meter*. The *device configurator* modifies the various configuration states of the device according the benchmarking requirements. The benchmarking process also loads the device with varying traffic patterns using the *traffic generator*. The *benchmark orchestrator* coordinates the various components in order to synchronize the configuration, workload and measurements from the power-meter. The collected information is then processed by an *analyzer* to generate various energy-proportionality indices and other power-related metrics.

A network switch or a router consists of many different components, such as the main chassis, linecards, TCAM, RAM, processor, fans etc. A complete instrumentation of various components is difficult to perform. First, we list various factors that are likely to affect power consumption for switches and routers; further, we describe the important ones that we have incorporated in the *device configurator*. While additional metrics and tests may have to be added to our suite in the future due to technological advances in switches and routers, using our current benchmarking suite, we are able to predict within a 2% error margin, the power consumed by a variety of switches deployed today in a real operational data center.

- *Base chassis power*: Higher-end switches (typically deployed at the edge of a network, and in data centers, etc.) come with a base chassis and a fixed number of slots. In each slot, a linecard can be plugged in. In lower-end switches (typically having 24 or fewer ports), the slots and linecards are fixed and cannot be changed. In both cases, the chassis power typically includes the power consumed by components in the switch such as processor, fans, memory, etc.

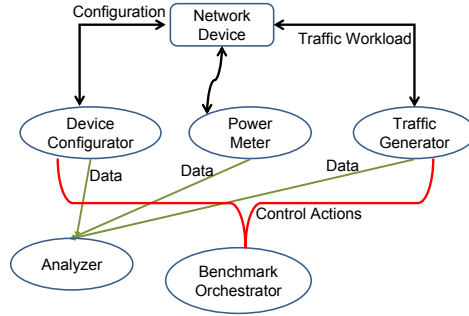


Fig. 1. Benchmarking suite architecture

- *Number of linecards*: In switches that support plugging in linecards, there is a limit on the number of ports per linecard as well as the aggregate bandwidth that each linecard can accommodate. This mechanism allows network operators the flexibility to only plug in as many linecards as they need. Further, it also offers choices such as the ability to plug in 24-port 1 Gbps linecard for an aggregate 24 Gbps capacity versus a 4-port 10 Gbps linecard for an aggregate of 40 Gbps capacity.
- *Number of active ports*: This term refers to the total number of ports on the switch (across all the linecards) that are active (with cables plugged in). The remaining ports on the switch are explicitly disabled using the switch’s command line interface.
- *Port capacity*: Setting this parameter limits the line rate forwarding capacity of individual ports. Typically, the capacity of a full-duplex 1 Gbps port can be also set to 10 Mbps and 100 Mbps.
- *Port utilization*: This term describes the actual throughput flowing through a port relative to its specified capacity. Thus, a port whose capacity has been set to 100 Mbps and having a 10 Mbps flowing through it has a port utilization of 10%.
- *Ternary Content Addressable Memory (TCAM)* : Switch vendors typically implement packet classification in hardware. TCAMs are supported by most vendors as they have very fast look-up times. However, they are notoriously power-hungry. The size of the TCAM in a switch is widely variable.
- *Firmware*: Vendors periodically release upgraded version of their switch/router firmware. Different versions of firmware may also impact the device power consumption.

The traffic characteristics at each port might also affect the power consumption; we list a few traffic factors that we have incorporated in the *traffic generator component* of our benchmark suite:

- *Packet size*: Varying from 48 bytes to 1500 bytes.
- *Inter-packet delay*: Time between successive packets at a port.

- *IP options set in the packet*: While this factor might not affect power consumption at switches performing MAC forwarding, processing packets that have IP options might impact the power consumption of a router.

Given these range of factors in the *traffic generator* and *device configurator* components of our benchmarking suite that can impact switch power consumption, our goal is to define the ones that will impact the switch power consumption the most. Capturing this short list of factors correctly, will allow us to be accurate in predicting switch power consumption.

3.2 Network Energy Proportionality Index

As noted previously, manufacturers need to ensure that devices consume energy proportional to their load. In the rest of this section, we describe how we can calculate energy proportionality for network devices.

In Figure 2, the power consumed by a device is plotted against the load on the device (in Gbps or active number of ports) with a maximum load of G Gbps. Ideally, the power consumed should be proportional to the load, with the maximum power consumed, M watts, being as low as possible. The curve marked *ideal* represents this desired energy proportionality behavior. In practice however, the behavior of network devices follow the line marked *measured*, with the device consuming I watts even under *idle* (no load) conditions. The difference between the *ideal* and *measured* lines forms the basis of the following *energy proportionality index (EPI)* for networking components, which we define as $EPI = \frac{M-I}{M} * 100$. If θ and ϕ are the angles at the origin for the ideal and measured power, EPI is simply $(\tan(\phi)/\tan(\theta)) * 100$. We express EPI in percentage, with 100 implying that the device has perfect energy proportionality and 0 implying that the energy consumed by the device is completely agnostic to offered load. Note that EPI is independent of the maximum load that can be carried by the device and thus is most useful in comparing the energy proportionality of devices in the same class with the same maximum load. *Normalized power* [8], another metric,




Fig. 2. Ideal (energy proportional) and measured power characteristics of network devices

is the maximum power consumed by the device divided by its aggregate bandwidth and is calculated as $NormalizedPower = \frac{M}{G}$. Since our devices span a wide range, from wireless access points to high end switches in data centers, we compute normalized power in milliWatts/Mbps. We use the above metrics to quantify the energy proportionality of a wide range of networking devices in the following section.

4 Experimental Details and Results

We build our benchmarking suite incorporating all the factors described in Section 3.1. To measure power consumed by a network device, we use a digital power-meter from Brand Electronics. In all of our experiments, we plug the power meter into an electrical outlet and the switch under test into the receptacle of the power meter. Our measurements include any power losses due to inefficiencies. We do not consider power factor in this paper. We use *expect* [9] scripts to build our *Orchestrator* and *Device Configurator* shown in Figure 1. Our configurator is modular so that we can easily add support for any other networking device as well. Our *Traffic Generator* consists of well-known programs such as iperf [10] as well as our own in-house packet generator program that sends CBR traffic with parameters to vary the packet size, inter-packet gap and total traffic throughput. As part of our benchmarking process, we run each experiment for 300 seconds and report the average power over the entire duration. Further, we run each experiment three times to minimize the effect of noise in these measurements. We observe insignificant differences between peak power and average power over the experiment duration.

As the first steps in our benchmarking process, for each network device, we measure the power consumed in the idle state (all ports inactive and no cables connected to them). We also measure the time taken for each switch to reach a steady state after booting. Table 1 lists the device categories we use in our study along with their maximum power ratings (plate power) as described in their published specifications. The measured maximum power is the power consumed when the offered load is equal to the total aggregate bandwidth the switch (or router) can support. We report all power measurements in Watts (except the last column). We anonymize the device models, but all of our experiments were performed on devices from well-known vendors such as Cisco, ProCurve, and Brocade.

Device A is a low-end network hub. While the maximum measured power for this switch is a relatively low 12.8 W, the fact that the switch has only 12 ports, with each port capable of forwarding traffic at a maximum of 100 Mbps, accounts for the rather high milliWatts / Mbps value when compared to devices B, C, D and E. Devices C, E and F can be built to different specifications with respect to the model and number of line cards that can be added to the chassis, while no such options exist for switches B and D. Switches B, C and E have the ability to support PoE (Power Over Ethernet), while D does not. Switch C is available as a modular chassis with 6-slots, with each slot capable of supporting

Table 1. Power consumption summary for network devices

Label	Type	Rated Max Power	Measured Max Power (M)	Measured Idle Power (I)	Time to reach steady state	EPI (in %)	Aggregate bandwidth in Mbps	mWatts / Mbps ^a
A	10/100 Hub	35	12.8	11.7	22 secs	8.59	1200	10.7
B	edge LAN switch	759 ^b	198	150	125 secs	24.2	48000	4.1
C	edge LAN switch	875 ^c	175	133.5	119 secs	23.7	48000	3.7
D	edge LAN switch	300	102	76.4	99 secs	25.1	48000	2.1
E	core switch	3000	656	555	212 secs	15.4	48000	13.7
F	edge router	300	210	168.5	195 secs	19.8	24000	8.75
G	wireless access point	12.5	8.3	5.2	50 secs	37.3	54	153.7

^a Measured max power in milliWatts / Aggregate bandwidth in Mbps. This term is equivalent to Joules per bit.

^b including 400W for PoE

^c including 400W for PoE

a 24-port linecard. Each port can be set to a maximum of 1 Gbps capacity. In our benchmarking experiments, we choose 2 linecards for a total of 48 ports. Switch E is a core switch, typically used as a root switch in data centers. This particular model comes with 9 slots. We plug in a single 48-port 1 Gbps linecard in one slot and the management blade in another slot for all our benchmarking tests. Device F is a router, designed to sit at or just before the network edge. It has 4 card slots, allowing us to plug in a 12-port linecard in each slot. Each port can operate at a line speed of 1 Gbps. We plug in 2 linecards for this device, for a total of 24 ports. Device G is a wireless access point with a capacity of 54 Mbps. While the actual power consumed by this device is the lowest compared to all other models, its efficiency, as described by milliWatts/ Mbps, is the lowest (153.7 W) due to the fact that the access point is only capable of supporting traffic at 54 Mbps.

Devices B, C and D have an EPI value around 25%, indicating that their energy consumption varies slightly with traffic, though not by very much. The core switch (E) and router (F) exhibit almost no energy proportionality. The wireless access point (device G) on the other hand has the highest EPI value. It is important to note that EPI does not translate into energy efficiency. In the case of G, the milliWatts that have to be used to forward 1 Mbps of traffic is the highest amongst all the devices; thus it is the least efficient device.

For a switch having a line card with 48 full-duplex 1 Gbps ports, one way to fully load the switch is to attach servers to each port and ensure 1 Gbps of traffic going in and coming out of each port or use expensive measurement equipment from vendors such as Ixia (www.ixiacom.com). We settle on a relatively inexpensive strategy to fully load the switch without using 48 sources. We loop-back 23 port pairs causing 46 ports to be active. For the remaining two ports in the switch, we attach a commodity linux server to each port.

We run our traffic generator program from one server; every packet from this server is sent to the broadcast IP address. This technique ensures that the traffic received by the switch port to which the server is connected is now forwarded to all ports on the switch, thus ensuring that all ports get the offered load. Since we have loops in our topology, enabling spanning tree and disabling spanning tree have separate effects.² While our technique of creating loops in a switch is unconventional³, it allows us to benchmark the switch power consumption with just a couple of servers instead of connecting a server to each switch port, which can be unscalable especially for larger switches.

For each active port on the switch (except model A and G), we measure the switch power consumption for 3 separate port capacities - 10 Mbps, 100 Mbps and 1 Gbps. For each capacity value, we vary the utilization of each port by changing the traffic flowing through it as well as by enabling and disabling spanning tree. As noted previously, we repeat our measurements for each configuration thrice, with each experiment lasting 5 minutes, for a total of 15 minutes.

To verify the validity of our looping technique for power measurements, we first benchmark switch D using the technique we describe above and note its power consumption for different switch and traffic configurations. Next, we validate our results by instrumenting the same switch with 48 different servers connected to each port, and for the same traffic and switch configurations. In each case, we find the results from our looping technique to be within 2 Watts of the 'attaching one server to each port' technique (the error margin was under 2%). While we could not validate all the other switches with this technique due to infrastructure constraints, based on conversations with switch vendors, we believe that our technique will hold for other switches as well.

We now describe results from our power benchmarking study with the edge LAN switch B. We obtain similar results for switches C, D and E, but don't report the results in detail here due to space constraints.

Edge LAN switch B: This 48-port switch, with a maximum per port capacity of 1 Gbps, is typically used in LANs as well as in data centers to connect servers in a rack. We find the power consumption to be very stable and the standard deviation and variance to be negligible. We plot average power consumed for each port capacity as a function of the number of active ports on the switch

² When spanning tree is enabled, the switch internally detects all the loops and ensures that for each port pair that form a loop, exactly one port is forwarding all the broadcast packets and the other port is receiving all the packets. For example, if all ports in the switch are enabled through the command line interface, and if one of the connected servers is generating 1 Gbps of traffic, each port in the switch either receives or forwards 1 Gbps of traffic. In the case when spanning tree is disabled, the loops in the topology cause the switch to flood every port with control packets. Every port in this case, both transmits and receives packets equal to its full line speed. In other words, if a port's capacity is configured to 1 Gbps, the throughput (due to the control packets) on this port is 1 Gbps in each transmitted and received state.

³ Communications with a switch vendor confirmed that our power measurement numbers using broadcast packets would match the observed power if we had connected a single server to each port and resorted to transmitting unicast packets.




Fig. 3. Switch B - 1 Gbps port capacity




Fig. 4. Switch B - 100 Mbps port capacity




Fig. 5. Switch B - 10 Mbps port capacity

in Figures 4, 4, and 5. Port capacity influences the power consumption significantly, especially for higher number of active ports. Further, as we increase the number of active ports, the impact of port utilization (whether no load or fully loaded) on power consumption is under 5%.

Moving onto the effects of traffic, we find that for a fixed traffic throughput at any port, packet size does not impact power consumption at all. In Table 2, we show the power consumed by the switch with 48 active ports as a function of 4 different packet sizes.

Next, we determine the effect of the number of entries in the TCAM on the power consumption. For this purpose, we add rules into the TCAM to accept and deny packets from certain IP addresses. We vary the number of TCAM entries from 0 to 3044, the maximum number available for this switch. With the TCAM

Table 2. Impact of packet size on power consumption

Number of active ports	Packet size in bytes	Average power in Watts
48	10	151.2
48	500	151.6
48	1000	151.4
48	1500	152.3

appropriately filled, we again repeat all our experiments for different switch and traffic configurations described above and find that the number of the TCAM entries does not impact power consumption at all. One interesting aspect that we observe is the effect of switch firmware on energy consumption. When we upgrade to the latest version of firmware available on the manufacturer’s website, and run our complete benchmarking suite again, we find that the power consumed is 8-12 W lower in every case. The idle power (I) is now 141 W, while the measured max power (M) is 178 W, for a new milliWatts/Mbps value of 3.7 and an EPI of 20.8%. We believe that concerns about energy usage is starting to motivate device manufacturers to make improvements in their firmware such as turning off unused components. While these improvements may be a contributing factor towards the lower power consumption, the EPI of the switch is still low.

We observe similar results for switches C, D, E and F. For devices where we can plug in linecards, we note that each linecard consumes 35-40 Watts. The base switch power now includes the chassis power plus the cost of each linecard. Other factors such as the number of active ports, effect of traffic, packet sizes, TCAMs, etc are similar to what we describe for switch B.

From the above results, we obtain a model to compute the total power consumed by the switch. From our measurements, we find that a linear model is able to accurately 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 use other models in that case.

Total power consumed consists of a fixed component (chassis power and power for each linecard) and a variable component that depends on the number of active ports, capacity of each port and utilization of each port.⁴ We experiment with a combination of port line rates; for example, we set the line speed of 16 ports in the switch to 10 Mbps, 16 other ports to 100 Mbps and the remaining 16 ports to 1 Gbps. We find that the switch power consumption follows our linear model even when individual ports have different line rates set. The power model is given by $Power_{switch} = Power_{chassis} + num_{linecards} * Power_{linecard} + \sum_{i=0}^{configs} numports_{configs_i} * Power_{configs_i} * utilizationFactor$. $Power_{linecard}$ is the power consumed by the linecard with no ports turned on, and $num_{linecards}$ is the actual number of cards plugged into the switch. Variable $configs$ in the summation is the number of configurations for the port line rate. $Power_{configs_i}$ is the power for a port running at line rate i , where i can be 10 Mbps, 100 Mbps, or 1 Gbps and $utilizationFactor$ is the scaling factor to account for the utilization of each port.

Validation: We use our linear power model to validate the power consumed by switches within a live data center in our labs. Most of the racks in our data center had two model D switches attached to all the servers in a rack. With help from network administrators, we were successful in connecting several model D switches to power meters to record the power as real traffic was flowing through

⁴ Our model is slightly different from the one proposed by [11] for routers; we also include the cost for each active port at its specified line rate and utilization.

them. Using SNMP queries, we obtain traffic flowing through each port for each switch, the number of active ports on the switch and the capacity of each port. We substitute the values from these live switches in our linear power model for switch D and find that our predicted power matches the real power measured by the power meter with an error margin of under 2%.

5 Related Work

Chabarek [11] *et al.* enumerate the power demands of two widely used Cisco routers; further the authors use mixed integer optimization techniques to determine the optimal configuration at each router in a sample network for a given traffic matrix. In our study, we consider a broader classification of switches and routers. Additionally, we also discuss energy proportionality and energy efficiency of these devices.

Some device manufacturers such as Cisco provide an online tool to calculate total power consumed by some of their switch and router models. However, these models are coarse-grained, and do not account for finer switch configurations such as the line rate at which each port has been configured as well as effects of traffic through each port.

At an individual switch or router level, researchers have been proposing techniques such as putting idle subcomponents (line cards, ports *etc.*) to sleep [1–4], as well as adapting link rates depending on the traffic [2–4, 12]. Allman *et al.* [13] suggest incorporating a power-aware proxy that relays keep-alive messages such as ARP replies to the switch on behalf of end devices such as desktops and laptops. This would allow the end devices to be put to sleep, while the proxy keeps the network connection alive.

Most of these proposals are in their early stages, and the underlying hardware support for implementing these schemes is not yet available, though equipment manufacturers are actively working on these proposals. While all these efforts are commendable and a step in the right direction, what has been lacking in the network community is understanding the energy savings by implementing these schemes. For example, how many Joules can we save by allowing a switch in the network edge to sleep for 30 secs? To the best of our knowledge, no in-depth measurement study exists that quantifies the actual energy consumed by a wide range of switches under widely varying traffic conditions as well the actual power savings that can be obtained by performing techniques such as link rate adaptation.

6 Implications and Conclusions

Benchmarking the power consumed for a variety of switch and router configurations is a challenging problem, that will continue to grow in importance. We build a benchmarking suite that incorporates several device configurations and traffic factors that can impact energy consumption and identify the main factors that impact power consumption the most. We summarize the main observations from our study:

- The power consumed depends on the number of active ports. Explicitly disabling unused ports on a line card reduces the device power consumption. As the number of active ports increases, power consumed increases linearly for all three port line speeds - 10 Mbps, 100 Mbps and 1 Gbps.
- The power consumed depends on the line speed each port is configured at, with 1 Gbps consuming the most power and 10 Mbps the least. This effect is due to the extra energy required to operate the physical layer (PHY) at higher line speeds. Thus, for ports that have low utilization, it is beneficial to configure the port line speed itself to 10 or 100 Mbps than imposing rate-limiting on traffic through the port without lowering the port speed.
- Traffic through the device (port utilization) does not have a significant effect on power consumed, as can be observed from the low EPI values of the devices. For 10 Mbps port line speed, the impact of traffic on power consumed is slightly higher than that of 1 Gbps line speed. Further, power consumption is independent of the packet size.
- Power consumed depends on other factors such as firmware version on the switch/router. In current models, number of TCAM slots used does not impact power. Making components such as TCAMs power-proportional will also lower the energy consumption of devices.

Our study has two implications for network operators and device manufacturers. Our measurements serve as motivation for device manufacturers to adopt more aggressive techniques such as turning off unused components in the devices, to make them more energy efficient. At the same time, the low EPI values of current devices suggest that techniques such as traffic consolidation might provide significant energy savings. For instance, network operators in an enterprise or data center networks might be able to consolidate traffic from multiple switches onto a single switch and turn off the used switches.

We have built a linear model for power consumed by network device; using this model, users can calculate the power consumed by their switch (or router) by specifying the configuration and traffic values for their device. Further, we define and compute the energy proportionality index for each switch, similar to what has been proposed for servers. We note that energy proportionality does not imply energy efficiency. These are two separate terms and energy management techniques should consider making devices both energy efficient as well as energy proportional. In the future, we plan to make the benchmarking suite available to build a comprehensive public repository of power characterization of all the network devices.

References

1. Gupta, M., Singh, S.: Greening of the Internet. In: Proceedings of ACM SIGCOMM (August 2003)
2. Gupta, M., Grover, S., Singh, S.: A feasibility study for power management in LAN switches. In: Proceedings of IEEE ICNP (October 2004)

3. Gupta, M., Singh, S.: Energy conservation with low power modes in Ethernet LAN environments. In: Proceedings of IEEE INFOCOM (MiniSymposium) (May 2007)
4. Nedeveschi, S., Popa, L., Iannaccone, G., Ratnasamy, S., Wetherall, D.: Reducing Network Energy Consumption via Rate-Adaptation and Sleeping. In: Proceedings of NSDI (April 2008)
5. Barroso, L.A., Holzle, U.: The case for energy-proportional computing. In: IEEE Computer (December 2007)
6. Mayo, R., Ranganathan, P.: Energy Consumption in Mobile Devices: Why Future Systems Need Requirements-Aware Energy Scale-Down. In: Falsafi, B., VijayKumar, T.N. (eds.) PACS 2003. LNCS, vol. 3164, pp. 26–40. Springer, Heidelberg (2005)
7. Rivoire, S., Shah, M., Ranganathan, P., Kozyrakis, C., Meza, J.: Modeling and Metrology Challenges for Enterprise Power Management. IEEE Computer (December 2007)
8. Al-Fares, M., Loukissas, A., Vahdat, A.: A Scalable, Commodity Data Center Architecture. In: Proceedings of SIGCOMM (August 2008)
9. Expect, <http://expect.nist.gov/>
10. iperf, <http://dast.nlanr.net/Projects/Iperf/>
11. Chabarek, J., Sommers, J., Barford, P., Estan, C., Tsiang, D., Wright, S.: Power awareness in network design and routing. In: Proceedings of INFOCOM (April 2008)
12. Gunaratne, C., Christensen, K., Nordman, B., Suen, S.: Reducing the Energy Consumption of Ethernet with Adaptive Link Rate (ALR). IEEE Transactions on Computers 57, 448–461 (2008)
13. Allman, M., Christensen, K., Nordman, B., Paxson, V.: Enabling an Energy-Efficient Future Internet Through Selectively Connected End Systems. In: Proceedings of HotNets (November 2007)