

University of Nebraska - Lincoln

DigitalCommons@University of Nebraska - Lincoln

CSE Conference and Workshop Papers

Computer Science and Engineering, Department of

2011

Rightsizing Bundle Link Capacities for Energy Savings in the Core Network

Lin Liu

University of Nebraska-Lincoln, linliu@cse.unl.edu

Byrav Ramamurthy

University of Nebraska at Lincoln, bramamurthy2@unl.edu

Follow this and additional works at: <http://digitalcommons.unl.edu/cseconfwork>



Part of the [Computer Sciences Commons](#)

Liu, Lin and Ramamurthy, Byrav, "Rightsizing Bundle Link Capacities for Energy Savings in the Core Network" (2011). *CSE Conference and Workshop Papers*. 194.

<http://digitalcommons.unl.edu/cseconfwork/194>

This Article is brought to you for free and open access by the Computer Science and Engineering, Department of at DigitalCommons@University of Nebraska - Lincoln. It has been accepted for inclusion in CSE Conference and Workshop Papers by an authorized administrator of DigitalCommons@University of Nebraska - Lincoln.

Rightsizing Bundle Link Capacities for Energy Savings in the Core Network

Lin Liu and Byrav Ramamurthy
Department of Computer Science and Engineering
University of Nebraska-Lincoln
Lincoln, Nebraska 68588-0115
Email: {linliu, byrav}@cse.unl.edu

Abstract—Current core networks are composed of high-end routers which are connected by high-speed fibers. These optical connections are commonly overprovisioned and in low utilization. Many of them are combined together to form bundle links or composite links and the component links are referred to as sublinks. These physical sublinks could be SONET connections, Ethernet circuits, wavelengths on a fiber, etc and they could be shut down or brought up independently. Selectively shutting down sublinks during low traffic periods could save a large amount of energy while keeping the network topology unchanged. Based on this concept, we propose a local heuristical threshold-based method to explore the potential energy-savings in the backbone network by adjusting the number of active sublinks in bundle links.

An experiment based on an Internet2 derived synthetic network was conducted to verify the performance of our method and the results show that 86% of energy consumed on ports of core routers could be saved when setting 90.0% as the link utilization threshold. The experiment also shows that setting 90.0% as threshold is safe enough to avoid data loss during extreme traffic increases in this case. Compared to previous proposed ILP (Integer linear programming) based global heuristic algorithms, our local heuristic algorithm can achieve energy-savings close to the optimum and greatly reduce the response time and the risk of data loss.

I. INTRODUCTION

Today's core networks carry much more voice and video data than ever before and the available network bandwidth has increased tremendously and this has led to the growth in network energy consumption. Higher network bandwidth necessitates more powerful processing capability and cooling systems, thus, consuming more energy. Powered wired networks in the United States alone costs an estimated 0.5-2.4 billion dollars per year [1]. Besides, current networks are designed to offer best-effort service and redundancy is provided everywhere. Overprovisioned link bandwidth and protection links are designed for peak traffic load or link failure which means that significant amount of energy is wasted. As a result, the link utilization of the network is pretty low (on average the link utilization of Internet Service Provider (ISP) network is estimated at 30% - 40%) [1]. Higher bandwidth also means greater variance of traffic [2]. Thus, network link bandwidth utilization can vary widely which also contributes to the low efficiency of network operation. Due to the above reasons, the telecommunications industry is listed as one of the least efficient industries in the world

and the situation would become worse in the future [3]. From the point of view of telecommunication carriers, energy consumption has more meaning than energy bills. Workloads exceeding the initial design capability of the energy system or the cooling system will force the ISP to upgrade or even physically reconstruct the whole site.

How to save energy in the core network has become an important area of networking research. Many methods have been proposed in this area from energy savings in network terminals such as personal computers or servers and even data centers to efficient operation of core routers such as designing routing protocols with power-awareness.

Our goal is to save energy in the ports of core routers with bundle links. When dealing with routers, an important issue arises which is different from dealing with network terminals, that is, whether or not to change the network topology while achieving network energy savings. As we know, shutting off a user computer or a server may not change the network topology and has very limited or no impact on the survivability or performance of the whole network; however, shutting down a core router or a bundle link attached to it, could have drastic consequences since it changes the network topology. From our understanding, changing network topology to achieve energy savings currently is not a good idea since a stable IP layer route is so critical that numerous upper-layer applications and protocols rely on it and frequent routing changes bring problems of packet loss, retransmission and serious delay, which are beyond the control of current routing protocols. Keeping network topology unchanged is a strict principle in our design of energy-efficient core networks which is different from some previous methods.

II. PROBLEM SETTING

Our method focuses on solving the energy problem under the scenario of a network with bundle links. Bundle link is not a new concept. Also termed as trunk aggregation, it has several aliases depending on the standard organizations or equipment providers, such as Ether-Channel (Cisco), Link Aggregation Control Protocol (LACP, IEEE 802.3ad), port channel, etc.

The formerly published bundle link standard is IEEE 802.3ad-2000 (LACP) which has now moved to IEEE 802.1AX standard and most applications conform to LACP. LACP is a layer 2 control protocol that can be used to

automatically detect, configure and manage multiple physical links between two adjacent LACP enabled devices into one bundle link.

Core routers of major vendors support LACP e.g. Cisco 10000 and 7600 series routers and Huawei Quidway S9300 Series Terabit Routing Switch. Cisco routers support bundle links across multiple chassis. The reasons to bundle multiple links together to work as one logical link are two-fold: bandwidth limitations and lack of resilience. In Ethernet and optical networks, the bandwidth upgrades are by an order of fixed multiple such as 10. For example, Ethernet upgrades from 1 Gbits/s to 10 Gbits/s. Before traffic demands reach the next scale, carriers prefer to manually configure to combine several Ethernet or SONET/SDH links into one logical link via link aggregation. Also, considering a single link connection, the cable itself or ports can fail and hence bundle link composed of multiple physical links could help to reduce the single points of failure. The average link utilization of sublinks of bundle link is normally below 40% [1]. This gives us a great opportunity to reduce network energy if we can dynamically adjust the number of active sublinks within one bundle link, that is, shutting down or bringing up sublinks according to current link traffic.

The capability of dynamically adjusting bandwidth between routers (and switches) is possible. One example is Huawei which has released the industry's first Hitless Adjustment of ODUflex (G.HAO) prototype, which enables efficient scaling of bandwidth to actual bandwidth demand, thus freeing up bandwidth resources and dramatically boosting transport efficiency.

There are two kinds of traffic distribution methods in core networks applying bundle link technique: bin-packing and load balancing. The former allows us to pack traffic into the minimum number of sublinks being used while the latter means that all the traffic is almost evenly distributed across all the sublinks. In the bin-packing case, the utilization of the first several sublinks and the utilization of the last sublink could be very different, while in load balancing, the sublinks have almost the same link utilization. In this paper, we only discuss the energy savings under bin-packing scenarios. The load-balancing case will be studied in our future work.

Another issue we need to consider is which mode we should shift the router ports to from the active status: low-rate mode or power down mode? The former is related to the rate scaling technique and the latter means totally shutting down the ports. Current research data shows that rate-scaling would not let energy consumption on the port decrease a lot while shutting the ports down could bring the energy consumption down to zero [4]. The time taken for a sublink to go from active status to down status could be on the order of milliseconds [5].

III. RELATED WORK

Gupta and Singh [6] suggest the idea of saving energy by putting the network interfaces and other router and switch components to sleep. They discussed the energy saving opportunities and how the sleeping mechanism works on the

network interface and accordingly impacts the current internet protocols. Chabarek [7] proposes a general model for router power consumption based on energy consumption measurement of different configurations of widely used core and edge routers. Along with mixed integer optimization techniques, they explore the potential impact of power-awareness in a set of example networks. Their experiments show power consumption can vary by as much as an order of magnitude. Nedeveschi [4] proposed two forms of power management: putting network components to sleep during idle times or adapting the rate of network operation to the offered workload. However these two kinds of methods need hardware support. Some other research studies focus on allowing the end systems to remain connected or disconnected [8]. Andrews et al. [5] study the routing and scheduling problem in a network in which each network element either works in zero-rate mode or full-rate mode, and the two schedules are compared when the routing is given as input.

Fisher et al. [1] investigated energy savings in bundle links for the first time. They formulated the problem of saving energy in bundle links as an NP-complete problem and they propose three ILP model based global optimization heuristic algorithms which are conducted in the network management system. They claim that the global optimization heuristic algorithms could achieve minimum number of live links, thus, achieving the maximum energy saving. And in their ILP formulation, they give an example in which k shortest paths could be shut down and traffic will be compressed into the last (longer) path to achieve great energy savings.

The method in [1] clearly requires that the network topology be changed. As discussed earlier, the consequence could be extremely complicated. The three global heuristic optimization algorithms they proposed (Fast Greedy Heuristic, Exhaustive Greedy Heuristic and Bi-level Greedy Heuristic) have a priori condition that the network administration systems collect bandwidth requirements of the whole network; but estimation of traffic demand could be difficult and imprecise especially in a large network. The ILP method is not a scalable solution. All these global heuristic algorithms take a long time to obtain the solution for ILP in a complicated network topology. For example for the Waxman topology, the fast algorithm took at least 30 minutes. We argue that the long processing time could be critical since during the period of half an hour, the traffic could exceed the capabilities of current active links. The chance of buffer overflow is pretty high. In conclusion the global heuristic ILP based methods are not scalable, are slow-acting, risky and come at the cost of a changing topology.

To solve these problems, we propose a local heuristic threshold based method to deal with energy saving in bundle links scenarios, which is local, autonomous, scalable, fast-acting and topology-invariant. Our method needs no involvement of the network management system, since the routers can make decision from their current link utilization information. Most importantly, our method keeps the network safe by keeping the network topology unchanged.

IV. PROBLEM STATEMENT

Internet2 is an U.S. research and education network which is designed to provide dynamic, innovative and cost-effective hybrid optical and packet network. In comparison to a commercial ISP network, Internet2 releases the information about its network topology and provides complete link traffic data in the form of RRD (Round Robin Database) file to download. The RRD tool fetches requisite data in a round robin way from multiple DS (data source) and consolidates the data in multiple methods and stores them in RRA (Round Robin Archives) as RRD files. Internet2 RRD files store last three years' link traffic data of Internet2 in a minimum interval of 10 seconds, and we consolidated them to intervals of 5 minutes for our experiment. With link traffic data of each five minutes extracted from RRD files, Internet2 link utilization can be calculated by standard formula. Our synthetic network utilizes almost the same network topology as Internet2 except that we assume bundle link in each link. Internet2 does not apply bundle link technique since a 10GE link is enough to carry its traffic; while in commercial ISP network bundle links are commonly applied. To simulate a large ISP network, in our synthetic network, we replace each link of Internet2 by a bundle link composed of ten 1GE optical sublinks, i.e., the same bandwidth but 10 times as many links than before. And we assume each sublink carries one-tenth of the link traffic of the original Internet2 link. In this way, the total traffic of each bundle link is the same as that of the original Internet2 link. The synthetic network topology is shown in Fig. 1 which includes 9 sites and 26 bundle links (totally 260 pair of 1GE bidirectional fibers). In the original Internet2 topology, between each two direct-connected nodes, there are two separate bidirectional links which serve different purposes and their link utilizations generally have big differences. We keep the same configuration in our synthetic network, thus, actually there are 26 bidirectional bundle links which exist in the topology. Table 1 shows information of all bundle links in our synthetic network. Thirty days (from July 1, 2010 to July 30, 2010) worth of link utilization data are extracted from Internet2 for experimental analysis and a typical link utilization data set is shown in Table 2. We start from the following observations: (1) A bundle link usually consists of several physical (SONET) links, which may have very low utilization. For example, bundle link 5423 from Atlanta to Houston consists of 10 physical links and their utilization is commonly less than 1%. (2) The energy consumed at the router ports due to a physical (SONET) link is related to its state (ON/OFF) and not related to its utilization. The method of rate adaptation does not contribute much energy saving here; ports must be shut down [4].

To avoid the problem of energy consumption differing in products from different vendors, we use port-hour or port-5minutes as energy-saving unit instead of watts for evaluation. Also note that, here we count the energy saving by the number of core router ports that could be shut down; while actually the energy savings are more than that since if the sublinks are




Fig. 1. Topology of Internet2

TABLE I
BUNDLE LINKS IN SYNTHETIC NETWORK

Site1	Site2	Bundle Link ID	Number of fibers
CHIC	ATLA	COMPOSITE_LINKCL5419	10
CHIC	ATLA	COMPOSITE_LINKCL5638	10
HOUS	ATLA	COMPOSITE_LINKCL5423	10
HOUS	ATLA	COMPOSITE_LINKCL5562	10
WASH	ATLA	COMPOSITE_LINKCL5133	10
WASH	ATLA	COMPOSITE_LINKCL5251	10
KANS	CHIC	COMPOSITE_LINKCL5567	10
KANS	CHIC	COMPOSITE_LINKCL5568	10
NEWY	CHIC	COMPOSITE_LINKCL5239	10
NEWY	CHIC	COMPOSITE_LINKCL5667	10
WASH	CHIC	COMPOSITE_LINKCL5250	10
WASH	CHIC	COMPOSITE_LINKCL5637	10
KANS	HOUS	COMPOSITE_LINKCL5560	10
KANS	HOUS	COMPOSITE_LINKCL5561	10
LOSA	HOUS	COMPOSITE_LINKCL5559	10
LOSA	HOUS	COMPOSITE_LINKCL5581	10
SALT	KANS	COMPOSITE_LINKCL5138	10
SALT	KANS	COMPOSITE_LINKCL5566	10
SALT	LOSA	COMPOSITE_LINKCL5563	10
SALT	LOSA	COMPOSITE_LINKCL5571	10
SEAT	LOSA	COMPOSITE_LINKCL5564	10
SEAT	LOSA	COMPOSITE_LINKCL5572	10
WASH	NEWY	COMPOSITE_LINKCL4643	10
WASH	NEWY	COMPOSITE_LINKCL5242	10
SEAT	SALT	COMPOSITE_LINKCL5565	10
SEAT	SALT	COMPOSITE_LINKCL5573	10

shut down, all the ports and fibers along that transmission path could be shutdown. Furthermore, if all the ports in a line card are shut down, the whole line card could be shut down and the energy savings of line card will be greater than that of shutting down its attached ports. To fully realize the energy savings in these scenarios might introduce additional complexity. To simply the calculation, we just count the number of ports shut down in core routers and prefer using percentage of ports shut down vs. total ports to evaluate the performance of our method.

The basic idea of saving energy in a bundle link is that during each time slot we carry the whole traffic using min-

TABLE II
TYPICAL LINK TRAFFIC INFORMATION EXTRACTED FROM INTERNET2
RRD FILE

Timestamp	Input	Output
1212296410	6.0812216354e-07	5.1619134628e-07
1212296420	5.0450804956e-07	4.5207642781e-07
1212296430	5.8053464978e-07	5.6639656430e-07
1212296440	5.6483092482e-07	5.5529660413e-07
1212296450	4.9683281349e-07	4.6534343965e-07
1212296460	5.2271588971e-07	4.9615828455e-07

imum number of sublinks within the bundle link, that is, when traffic is low, we shut down sublinks and shut down the associated ports of routers, while retaining the Internet2 layer3 adjacencies; when traffic rises high enough, we bring up new sublinks to satisfy the higher bandwidth requirement. As discussed earlier, at any time, each bundle link should keep at least one sublink active. The problem is to determine the minimum number of active sublinks and to determine the number of sublinks which should be shut down or brought up. A global heuristic optimization method [1] was previously proposed which however could result in some problems. The first problem is that it needs knowledge of traffic requirements of the whole network which is hard to estimate. The second problem is that it is an ILP (Integer Linear programming) model based solution and solving this ILP could be time consuming depending on the size and complexity of network, which makes this method not scalable. The disadvantages of this method were fully discussed in the related work section above.

We propose a local heuristic threshold based method which performs decision making at local core routers and sets utilization thresholds to trigger signals of shutting down and bringing up the sublinks. The basis of the method is that a router has the best knowledge of its current link traffic and setting appropriate sampling frequency and utilization threshold of link traffic could make the router sensitive and agile enough to respond to the varying traffic. In our method, the utilization threshold is set as the major parameter to decide the number of active sublinks for each bundle link. Only when the utilization of latest activated sublink goes beyond the threshold, does the router bring up a sublink to allow for increased traffic and when link utilization of latest activated sublink reaches zero, the sublink is shutdown to save energy. Thus there are two thresholds which are used by a core router. A high threshold (t_{high}) which triggers adding a sublink and a low threshold (t_{low}) which triggers shutting down a sublink. In our study, we always set t_{low} to be zero, so that we do not accidentally lose any traffic. However, we set t_{high} threshold to a high value (such as 90%) to study its impact on the energy savings. In the remainder of the paper, we refer to the t_{high} threshold as simply threshold.

V. OUR APPROACH

First, we would like to estimate the potential energy savings in energy (port-hours) by efficiently carrying traffic over the minimum number of sublinks for each bundle link at each different time slot (one time slot is five minutes) during the day.

For each bundle link, its number of active sublinks is decided independently and locally. Ideally, the traffic demand of bundle links in each time slot are known in advance. Let us say for one bundle link L, the traffic demand in the next time slot is D. The number of sublinks in it is N and currently the number of active sublinks is C. Each sublink has bandwidth B and utilization threshold of L is set as T. Then the formula for sublink determination for next time slot is as follows:

$$A = \lceil \frac{D}{(B * T) * N} - C \rceil \quad (1)$$

If A is positive, it means that A sublinks need to be brought up; zero means nothing to do; otherwise, A sublinks could be shut down. One problem with this method is that D is hard to predict precisely so it is expected that there is variance between the estimated traffic and the real traffic. So Equation 1 can be applied to calculate the optimum value but is not practical. Actually since we do not know how much traffic increases or decreases in the next five minutes, we can only speculate on the traffic demand of the next time slot based on the actual traffic in previous time slots. And a simple way to operate the network is by setting an alarm line just like a flood alarm line, and only if the traffic rises to or beyond the line, the alarm signal will be triggered. Better traffic prediction methods based on some statistical models could be developed to improve the performance. Also in practice service providers care about the link utilization more than the amount of link traffic passing through. So the first step of our calculation is translating the link traffic into the link utilization. Then some utilization headroom should be reserved for the bursty traffic, and hence, a special link utilization threshold should be set up. When the link utilization of the last sublink in the previous time slot reaches to or beyond the utilization threshold, the signal to add sublinks would be triggered, and in the next time slot, one or more new sublinks will be activated to adapt the rising traffic. The number of newly added sublinks can be one, or two or more depending on the character of link traffic and how aggressive or conservative the strategy that is adopted. If the traffic is always bursty, each “add-sublink” signal should activate more sublinks to allow all of the rising traffic in order to reduce the possibility of overflow; otherwise, adding one more sublink each time might be enough. In the specific Internet2 case, since some links are reserved for education or research purposes, their link loads turn out to be more bursty and unpredictable than those of normal ISP links. In our experiments, we checked cases of adding one and two new sublinks for comparison of the possibilities of overflow.

Assigning an appropriate threshold value for each bundle link is also tricky. Firstly the threshold cannot be too high

(very close to 100%), since bringing up a sublink takes time (on the order of milliseconds) and the new sublink may not be ready to carry traffic and a large quantity of data might be lost. Secondly it is inefficient if the threshold is set too low (such as 40%). A tradeoff is made between the two extreme scenarios to keep the system sensitive and efficient. In our experiment, a preliminary threshold was set as 90% for each bundle link to check the energy saving and risk since 90% is close to the link capacity while still having some headroom to adapt to the bursty traffic. Also we conduct the same experiment but set the threshold as 80% for comparison. An optimal threshold value could be calculated from statistical analysis of long-term historical utilization data.

In addition to setting the threshold to reasonable values (such as 90%), we set the threshold to 100% and assume that the link traffic demand of each time slot is known in advance in order to calculate the theoretical optimum (least) number of sublinks needed for each bundle link to carry all of its traffic in each timeslot. Note that as pointed out earlier, this optimum (maximum savings) cannot be reached in practical settings due to high traffic variation and risk of data loss. We explain this computation in more detail below.

VI. EXPERIMENTS AND DISCUSSION

Firstly, we would like to know the maximum value of potential energy saving in our synthetic network within 30 days, that is, the upper bound of our local heuristic method. RRD files provide link traffic data in each time slot (5 minutes). From that the bundle link traffic at each time slot could be calculated. With the threshold set as 100%, the number of active sublinks needed for each bundle link at each individual time slot can be calculated and the number of sublinks shut down are easily figured out. Summing up the inactive ports for all the time slots and all the bundle links, the energy saving in a synthetic network for each day, each week, and the whole 30 days could be calculated and that is the upper bound. This is shown in the figures (Fig. 2, Fig. 3, Fig. 4) as dotted lines.

Secondly, an experiment using local heuristic threshold-based algorithm was conducted to check the performance. In the experiment, the threshold is set as 90% for each bundle link. We would like to know whether setting the threshold as 90% is safe enough or not to handle the varying traffic in this scenario, that is, we wish to verify whether one more active sublink is enough or not to carry the new traffic.

There could be extreme traffic variations that can lead the bin-packing threshold algorithm to produce traffic overflow. Imagine that the traffic increases too much in the next time slot, thus, one new sublink is not enough to carry the new traffic, and data could be lost. Add two more active sublinks each time new sublinks is needed can help reduce the possibilities of traffic overflow at the cost of slightly decreased energy savings. Fig. 2 shows the energy saving of the last time slot (23:55:00) for 30 days. Note that the unit of y axis is port*5minutes. The energy saving in one day is shown in Fig. 3. The energy saving in 30 days is shown in Fig. 4.




Fig. 2. Energy savings of last time slot (23:55:00) for 30 days setting 90% as the threshold.




Fig. 3. Energy savings of all bundle links in one day (20100701) setting 90% as the threshold.




Fig. 4. Energy savings of all bundle links over 30 days (20100701-20100730) setting 90% as the threshold.

During almost the whole month of July 2010, all bundle links in the synthetic network can save around 86% of total ports energy which is more than 160,000 port*hours in total.

Since the link traffic of the next time slot is unforeseeable, the simplest and riskiest way of local heuristic method is always to add one or two active sublinks for the next time slot if the utilization threshold is exceeded, since we believe 5 minutes is small enough time slot. Based on this assumption, we executed the local heuristic threshold method on our synthetic network for a duration of 30 days (July 1, 2010 to July 30, 2010) and obtained some interesting preliminary results.

The result of applying a local heuristic method shows that setting 90% as the threshold and always adding only one active sublink in the next time slot for each bundle link if it is needed, traffic overflow occurs in less than 6111 of the 224,640 timeslots in the 30 days period across all the bundle links in the network (or 2.7% of the timeslots); while changing the setting to always add two active sublinks each time the utilization threshold is exceeded, the timeslots of traffic overflow reduced to 0.18%. Even in the latter setting, the energy savings is still considerable. The above experiment shows that in our synthetic network, with appropriate setting, 90% is a safe enough utilization threshold to withstand the extreme traffic case. Fig. 2 shows that in a single timeslot (the last timeslot of a day) over 30 days, energy of 188 port*5minutes could be saved. Fig. 3 shows that on a particular day, energy of 63522 port*5minutes could be saved. Fig. 4 shows that on a particular 30 day period, energy of 1,926,650 port*5minutes could be saved. We can see that in all three cases most ports can be shut down and significant amounts of energy can be saved. Results from additional experiments (with 80% as threshold) provide similar but decreasing savings in energy consumption.

VII. CONCLUSION AND FUTURE WORK

Observations from link traffic data of Internet2 show that substantial energy savings could be achieved in bundle links of the core network. A preliminary experiment was conducted to explore the upper bound of energy savings. Setting 90% as threshold to evaluate the performance of our method of “rightsizing bundle links” resulted in considerable energy savings with very small risk of data loss.

The experimental results show that although traffic is unpredictable, with appropriate setting of bundle link threshold, our method works well and achieves energy savings close to

optimal without substantial data loss. Ignoring the cost of port operations (shutting down and bringing up), the energy savings achievable is very promising, i.e., 86% energy can be saved. Due to local decision making, our method responds quickly and is more effective than global heuristic methods.

Several tasks are left as future work such as on-line model evaluation and deployment of a Round-Robin algorithm for router port allocation. A protocol to synchronize the steps for two routers’ port operations could be developed. Also, more complicated bundle link utilizations setting for the load balancing case can be researched. More precise calculation by using the shorter time interval and the longer period which is more than 30 days could be conducted in the future. Finally, the alarm mechanism in network management needs to be modified to deal with the issues related to port operations of rightsizing bundle links.

VIII. ACKNOWLEDGEMENT

This work was supported in part by DOE (DE-SC0001277) and NSF (CNS-1040765).

REFERENCES

- [1] W. Fisher, M. Suchara, and J. Rexford, “Greening backbone networks: reducing energy consumption by shutting off cables in bundled links,” in *Proceedings of the first ACM SIGCOMM workshop on Green networking*, ACM, 2010, pp. 29–34.
- [2] R. Morris and D. Lin, “Variance of aggregated web traffic,” in *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 1. IEEE, 2000, pp. 360–366.
- [3] D. Seo, A. Featherstone, D. Weisman, and Y. Gao, “Market consolidation and productivity growth in U.S. wireline telecommunications: Stochastic frontier analysis vs. malmquist index,” *Review of Industrial Organization*, vol. 36, no. 3, pp. 271–294, 2010.
- [4] S. Nedeveschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall, “Reducing network energy consumption via sleeping and rate-adaptation,” in *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*. USENIX Association, 2008, pp. 323–336.
- [5] M. Andrews, A. Anta, L. Zhang, and W. Zhao, “Routing and Scheduling for Energy and Delay Mimimization in the Powerdown Model,” in *IEEE INFOCOM*, 2010.
- [6] M. Gupta and S. Singh, “Greening of the Internet,” in *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications*. ACM, 2003.
- [7] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, “Power awareness in network design and routing,” in *IEEE INFOCOM*, 2008.
- [8] M. Allman, K. Christensen, B. Nordman, and V. Paxson, “Enabling an energy-efficient future internet through selectively connected end systems,” in *6th ACM Workshop on Hot Topics in Networks (HotNets)*, 2007, pp. 1–6.