PROPOSAL TITLE: Green Networking: Reducing the Energy Use of LAN Switches and Connected Hosts

---

PRIMARY INVESTIGATOR

Ken Christensen (Professor)
Department of Computer Science and Engineering
4202 East Fowler Avenue, ENB 118
University of South Florida
Tampa, FL 33620
Email: christen@csee.usf.edu

PROJECT SUMMARY

Reducing the energy consumption of LAN switches and routers, and of connected hosts, is of interest for environmental, economic, and operational reasons. In this project, we propose to investigate 1) the ability to cycle an Ethernet link between on and off using PAUSE flow control and powering down during the off periods, and 2) using existing Cisco switch and router deep packet inspection capabilities to control traffic to sleeping hosts enabling the hosts to sleep longer and not be woken-up by spurious traffic. Our proposed Ethernet PAUSE Power Cycle (PPC) can allow a switch core or line card to largely power down during coordinated link off times. Preliminary results show that a 100 Mb/s edge Ethernet link with a 10 millisecond 50% on/off cycle has little, or no, perceivable effect to a typical user. PPC can be compatible with *existing* Ethernet NICs (unlike the proposed IEEE 802.3az Energy Efficient Ethernet Rapid PHY Selection method). Investigation of how to best schedule on/off cycles is needed. Internet hosts can sleep when inactive, but need support from the network to control spurious traffic that can cause unnecessary wake-ups. We will investigate how host power state can be communicated to a switch, and how existing capabilities (such as Cisco IPS and/or PISA application awareness) can be used to "protect" a host while it sleeps and still maintain full network presence of the host. Requested are $80,000 and equipment donations to support two Ph.D. students for two years. The students will be available for summer internships in 2008 and 2009 making it possible to directly exchange knowledge and research results with Cisco.

DESCRIPTION OF RESEARCH AND GOALS

The economic and environmental impact of electricity usage by IT equipment continues to grow. Energy use is one of the major cost factors in provisioning and operating data centers. The energy use of all office and network equipment has been estimated to be 74 TWh/yr in 2000 [1]. The energy consumption of hubs, switches, and routers in the Internet was estimated to be about 6 TWh/yr in 2000 and is expected to increase significantly in the near term [2]. Several consortium-level initiatives are now underway to reduce energy use of IT equipment [3, 4]. Standards initiatives for reducing energy used by networks include IEEE 802.3az Energy Efficient Ethernet [5]. In the commercial arena, D-Link is now offering "Green Ethernet" switches based on variable link power as a function of link length and turning-off unused ports [6]. The goal is for IT equipment – including networking equipment – to achieve power use proportional to utilization. In current IT equipment, power use mostly a function of capacity and not utilization. Figure 1 shows typical (actual) and linear (desired) power use as a function of utilization. This project addresses two ideas for moving LAN switches and hosts towards a more linear power use profile.

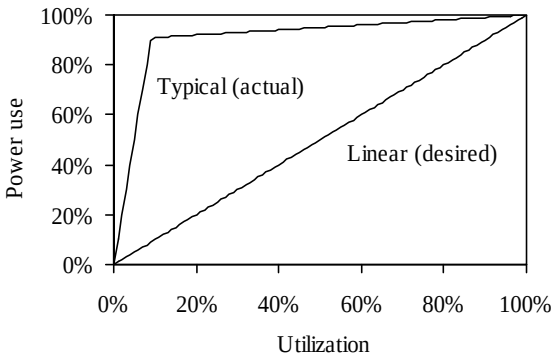

Figure 1. Utilization versus power use for IT equipment

Our goal is to explore two distinct ideas:

1) How PAUSE flow control can be used to cycle Ethernet links between on and off such that during coordinated off periods internal hardware in a switch could be powered-down to save energy. Key to this problem is understanding the effects of link off periods on higher layer protocols and network applications, and how to best schedule such link off periods to minimize performance impact to users while maximizing energy savings.

2) How Cisco switch and router deep packet inspection capabilities, including Cisco IPS capabilities [7] and PISA [8], can be used to control traffic sent to sleeping hosts. Such traffic filtering can minimize spurious wake-ups and enable greater energy savings in connected hosts. Key to this problem is finding a way for hosts to communicate their awake/sleep state to a switch and to determine how/what traffic the switch should filter to maximize host sleep time while still allowing access by legitimate remote applications.

We will develop a test bed and conduct simulation and experimental evaluations for both research directions.

Research Plan for Idea #1 – Pause Power Cycle (PPC) for Reducing LAN Switch Energy Use

Ethernet links are very lightly utilized – in general [9], traffic is bursty with infrequent high rate peaks. The proposed PAUSE Power Cycle (PPC) exploits this property to reduce energy use [10]. The basic idea of PPC is to use existing Ethernet PAUSE flow control frames sent from a first-level LAN switch to connected hosts (and other network equipment) to stop traffic flow on links incoming to the switch. During a coordinated link off time, switch components such as the switch core and/or line cards can be powered down. Figure 2 shows a switch with PPC functionality sending PAUSE frames to connected hosts and the edge router. Obviously, a major concern with PPC is what effects it may have on higher layer protocols and application performance. PAUSE frames sent to hosts may cause applications to block (but, there will be no packet loss). PAUSE frames sent to other network equipment may cause packet loss due to buffer overflow. These effects will only occur if a link off period coincides with a traffic burst. If packet loss does occur, TCP (as used by most network applications) will likely recover the lost packets, possibly even invisibly to a user.

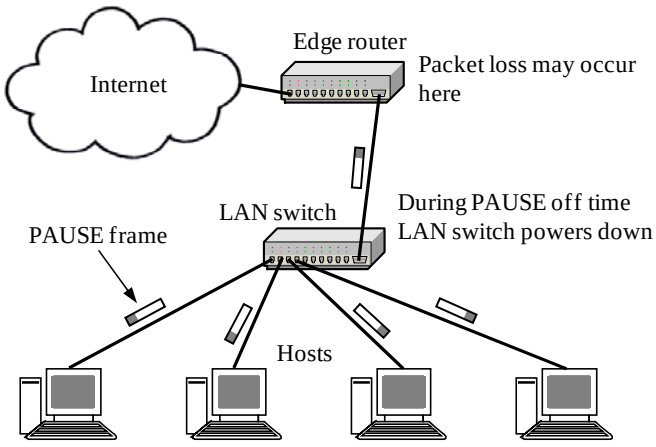

Figure 2. PAUSE Power Cycle (PPC) in a LAN switch

We will conduct an experimental and simulation based performance evaluation of PPC. The performance evaluation will entail greatly expanding a PPC emulation test bed that we have already constructed [10] and building an ns2 model of PPC. Our current PPC emulation experiments show that with cycles of 10 milliseconds on and 10 milliseconds off for a 100 Mb/s link, test subjects cannot detect a difference over a full bandwidth link. In our preliminary experiments, test subjects were web surfing, emailing, watching YouTube videos, and talking on Skype [10]. These results may not be surprising given that very often web users can be just as “happy” on a shared and low bandwidth WiFi wireless network as on a high-bandwidth dedicated link of a wired network.

PPC builds on work by the IEEE 802.3az Energy Efficient Ethernet task force [5]. PPC can be compatible with existing Ethernet NICs (unlike the proposed IEEE 802.3az Rapid PHY Selection methods). The outcomes of our performance evaluation will identify the cases where PPC can be employed and will seek to develop an adaptive PPC that can adjust to user patterns. Clearly, PPC is not suitable for all Ethernet links – but, we believe it may be suitable for many Ethernet links most of the time. We will also develop a software switch (a PC with multiple NICs) as a test fixture and/or modify a real switch with Cisco assistance to actually implement PPC.

Research Plan for Idea #2 – Switch Support for Extending Host Sleep Time

Significant amounts of energy are consumed by inactive network hosts forced, or induced, to remain fully powered on by network protocols and applications [1]. For example, to maintain reachability a host must reply to ARP packets. A host can be set to wake-up on ARP packets. It would in this case wake-up very often given that ARP packets flow very frequently on most links. We have measured a rate of over 10 directed ARPs per hour (these are all ARPs destined to the host) to an idle desktop PC connected to the university network. Other packet flows – at protocol and application layers – also require attention by a host. Network proxying has been explored as a means to relieve a host of having to be fully powered-on at all times to deal with relatively light weight, network-related duties [11, 12, 13]. We propose to investigate how a LAN switch or router can serve to "protect" a sleeping host to allow it to sleep for an extended time and still allow it to maintain network presence. There are four key aspects to this problem:

1) How can a host inform a switch/router of its power management state (awake, going to sleep, sleeping, waking-up)?

2) What packets should be filtered-out or otherwise controlled at a switch/router and how can this be done using existing capabilities?

3) For what packets and/or events should a switch/router wake-up a sleeping host and how can it best do so?

4) What protocols might a switch/router be able to execute on behalf of a sleeping host? For example, can a switch/router respond to ARP packets for a sleeping host? Can it maintain a DHCP leased IP address?

Beyond simple lower-layer protocol flows are application flows, often within a TCP connection. For example, IM uses application level keep-alive flows to maintain an IM connection to an IM server [14]. Is it possible for a switch to be part of such flows (e.g., as a "man in the middle" of a TCP connection) and enable hosts to sleep while being woken-up only for user specified interactions?

Cisco routers and switch products increasingly include deep packet inspection capabilities. For example, IPS capabilities are offered in 87x series of routers [7] and supervisor engines in Catalyst switches can control traffic flows (such as with PISA in some Catalyst switches [8]).We plan to develop prototype packet filters and proxy capabilities in a software switch (keeping in mind that these capabilities must be portable to Cisco products). We will study how a host can inform a switch of its power management state. Can there be implicit detection, or must there be explicit notification, of host power state by a switch? Clearly, there are many open questions, and we will address them with simulation modeling, system building and experimentation. Our experimentation will include measurement of potential energy savings as the extended time that a host can sleep.

We will also monitor and explore how current directions in network connectivity proxying [15] (which may be required by future EPA Energy Star specifications for PCs [16]) could be implemented in switches. Such an implementation could offer a competitive advantage for Cisco.

Previous and Related Work

Researchers at Lawrence Berkeley National Laboratory (LBNL) and the principal investigator (Christensen) at the University of South Florida were among the first to consider how energy consumption of network equipment and network-connected equipment could be reduced. A "Green TCP" and proxying of ARPs and low layer protocols were first investigated in the mid-1990s [11, 17]. In 2003 two seminal papers appeared (one from the principal investigator and the other from Suresh Singh at Portland State University) to describe the challenges of reducing energy use of networks [18, 19]. Christensen and his student explored Adaptive Link Rate (ALR) for Ethernet [20, 21] to reduce direct energy use of network links and proxying [12, 20] to reduce induced energy use of network connected hosts. The work on ALR directly fostered the IEEE 802.3az Energy Efficient Ethernet effort [5].

Singh and his students have explored power management for LAN switches [22, 23, 24]. They have investigated Ethernet traffic and how there may be sufficient idle time between packets to power down internal switch components. The proposed work on PPC is inspired by both the previous work of the principal investigator in ALR

and the work by Singh and his students. A very recent work by Christensen and colleagues explored wide ranging architectural primitives for supporting selective connectivity of networks hosts. Selective connectivity is aimed at improving energy efficiency of network hosts by allowing network hosts to sleep and maintain full network presence [25]. This proposed project builds on previous and related work in ways that will be of direct benefit to Cisco.

- -

REFERENCES

[1] K. Kawamoto, J. Koomey, B. Nordman, R. Brown, M. Piette, M. Ting, and A. Meier, "Electricity Used by Office Equipment and Network Equipment in the U.S.: Detailed Report and Appendices," Technical Report LBNL-45917, Energy Analysis Department, Lawrence Berkeley National Laboratory, February 2001.

[2] K. Roth, F. Goldstein, and J. Kleinman, "Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings Volume I: Energy Consumption Baseline", National Technical Information Service (NTIS), U.S. Department of Commerce, Springfield, VA 22161, NTIS Number: PB2002-101438.

[3] The Green Grid: Improving Power Efficiency in Data Centers, 2007. URL: http://www.thegreengrid.org/home.

[4] Climate Savers Computing, 2007. URL: http://www.climatesaverscomputing.org/.

[5] IEEE 802.3 Energy Efficient Ethernet Study Group, 2007. URL: http://grouper.ieee.org/groups/802/3/eee_study/index.html.

[6] D-Link Innovative Green Ethernet, 2007. URL: http://www.dlink.com/corporate/environment/dlink-green-ethernet/.

[7] "Cisco IOS IPS Value-Added Security Services in Cisco IOS Software", 2007. URL: http://www.cisco.com/application/pdf/en/us/guest/products/ps6540/c1031/cdccont_0900aecd80644585.pdf.

[8] "Application Intelligence and Integrated Security Using Cisco Catalyst 6500 Supervisor Engine 32 PISA," 2007. URL: http://www.cisco.com/en/US/products/hw/modules/ps2797/prod_brochure0900aecd805aeb81.html.

[9] A. Odlyzko, "Data networks are lightly utilized, and will stay that way," Review of Network Economics, Vol. 2, No. 3, pp 210-237, September 2003.

[10] F. Blanquicet, "PAUSE Power Cycle; A New Backwards Compatible Method to Reduce the Energy Consumption of Ethernet Switches," to be submitted to the *Ethernet Alliance White Paper Challenge*, 2008. URL: http://www.ethernetalliance.org/join/eaup/whitepaper_challenge/.

[11] K. Christensen and F. Gulledge, "Enabling Power Management for Network-Attached Computers," *International Journal of Network Management*, Vol. 8, No. 2, pp. 120-130, March-April 1998.

[12] M. Jimeno and K. Christensen, "A Prototype Power Management Proxy for Gnutella Peer-to-Peer File Sharing," *Proceedings of the IEEE Conference on Local Computer Networks*, pp. 210-212, October 2007.

[13] Low Power V 1.0, Universal Plug and Play Forum, August 28, 2007. URL: http://www.upnp.org/specs/lp.asp

[14] R. Jennings III, E. Nahum, D. Olshefski, D. Saha, Z-Y. Shae, and C. Waters, "A Study of Internet Instant Messaging and Chat Protocols," *IEEE Network*, Vol. 20, No. 4, pp. 16-21, July/August 2006.

[15] B. Nordman and K. Christensen, "Improving the Energy Efficiency of Ethernet-Connected Devices: A Proposal for Proxying," Version 1.0, *Ethernet Alliance White Paper*, September 2007. URL: http://www.ethernetalliance.org/technology/white_papers/Proposal_for_Proxying_edit.pdf.

[16] U.S. EPA, "ENERGY STAR Program Requirements for Computers", September 2006. URL: http://www.energystar.gov/index.cfm?c=archives.computer_spec.

[17] L. Irish and K. Christensen, "A 'Green TCP/IP' to Reduce Electricity Consumed by Computers," *Proceedings of IEEE Southeastcon*, pp. 302-305, April 1998.

[18] M. Gupta and S. Singh, "Greening of the Internet," *Proceedings of ACM SIGCOMM*, pp. 19-26, August 2003.

[19] K. Christensen, "The Next Frontier for Communications Networks: Power Management," *Proceedings of SPIE Performance and Control of Next-Generation Communications Networks* (extended abstract), Vol. 5244, pp. 1-4, September 2003.

[20] C. Gunaratne, K. Christensen, and B. Nordman, "Managing Energy Consumption Costs in Desktop PCs and LAN Switches with Proxying, Split TCP Connections, and Scaling of Link Speed," *International Journal of Network Management*, Vol. 15, No. 5, pp. 297-310, September/October 2005.

[21] C. Gunaratne, K. Christensen, S. Suen, and B. Nordman, "Reducing the Energy Consumption of Ethernet with an Adaptive Link Rate (ALR)," to appear in *IEEE Transactions on Computers*.

[22] M. Gupta, S. Grover, and S. Singh, "A Feasibility Study for Power Management in LAN Switches," *Proceedings of the 12th IEEE International Conference on Network Protocols*, pp. 361-371, October 2004.

[23] M. Gupta and S. Singh, "Energy Conservation with Low Power Modes in Ethernet LAN Environments," *Proceedings of IEEE INFOCOM*, May 2007.

[24] M. Gupta and S. Singh, "Dynamic Ethernet Link Shutdown for Power Conservation on Ethernet Links", *Proceedings of the IEEE International Conference on Communications*, pp. 6156-6161, June 2007.

[25] M. Allman , K. Christensen, B. Nordman, and V. Paxson, "Enabling an Energy-Efficient Future Internet Through Selectively Connected End Systems," *Proceedings of the Sixth Workshop on Hot Topics in Networks*, November 2007.

TIME FRAME FOR FUNDING AND RESEARCH COMPLETION

It is requested that funding begin on August 1, 2008 and ends on May 31, 2010 for a total of four academic semesters. This project will serve as a partial continuation of a currently NSF funded project in the area of reducing energy use of network edge devices. See http://www.csee.usf.edu/~christen/energy/main.html.

Research Milestones:

Summer 08	Summer internships at Cisco for students (Miguel Jimeno and Francisco Blanquicet)
1 Aug 08	Project starts (as continuation of NSF funded work)
15 Dec 08	Completion of 1) PPC simulation model and 2) traffic control for sleeping in a PC testbed
31 May 09	Demonstrate emulated PPC in a test bed and simulated PPC in an ns2 simulation and demonstrate use of packet filtering to enable greater host sleeping time
Summer 09	Summer internships at Cisco for students (Miguel Jimeno and Francisco Blanquicet)
1 Aug 09	Project starts (as continuation of NSF funded work)
15 Dec 09	Demonstrate PPC in a software (or actual modified) switch and demonstrate actual packet filtering and control for extending sleep time of hosts
31 May 10	Completion of technology transfer to Cisco and publication in conferences and journals

---

ANY REQUIRED/EXPECTED RESEARCH COOPERATION WITH CISCO

Andrea Baldini is the internal Cisco Champion for this project. Cooperation could include summer internships in 2008 and 2009 for the funded Ph.D. students.

SHORT BIOGRAPHIES OF THE RESEARCHERS

Ken Christensen (primary investigator – university faculty)

Ken Christensen received his Ph.D. in Electrical and Computer Engineering from North Carolina State University in 1991. He is a Professor at the University of South Florida. Previous to joining the University of South Florida, he was an engineer at IBM/RTP for ten years. His research and teaching interest is in performance evaluation of computer networks. He has over sixty conference and journal publications and thirteen U.S. patents (including patents in the area of LAN switching such as adaptive cut-through used on early IBM and Kalpana Ethernet switches). In 1998 and 1999 he was awarded a NASA/ASEE summer faculty fellowship at Kennedy Space Center. In 1999 he was awarded a CAREER grant from the National Science Foundation. Ken is a licensed Professional Engineer in the state of Florida, a member of ACM and ASEE, and a senior member of IEEE. Ken will be the editor-in-chief for the *International Journal of Network Management* starting in January 2008 and is on the editorial boards for *Computer Communications* journal and *IEEE ITPro* magazine. His homepage is at http://www.csee.usf.edu/~christen.

Francisco Blanquicet (funded graduate student)

Francisco Blanquicet is a Ph.D. student in the Department of Computer Science and Engineering at the University of South Florida. He graduated with a B.S. in Computer Engineering and a B.S. in Computer Science from the University of South Florida in 2006. His research interests are in performance evaluation of computer networks and dynamic power management of network devices. Presently, he is investigating methods to reduce the energy consumption of Ethernet switches and studying how existing protocols can be extended for communicating power state in networks. His expected graduation date is December 2010 and he will be seeking an industry position upon graduation. His homepage is at http://www.cse.usf.edu/~fblanqui.

Miguel Jimeno (funded graduate student)

Miguel Jimeno is a Ph.D. student in the Department of Computer Science and Engineering at the University of South Florida. He graduated with a B.S. in Computer Engineering from the Universidad del Norte (Colombia) in September 2002. He is a staff member on leave from the same university where he will return after his Ph.D. graduation. He has been studying at the University of South Florida since Fall 2005. His research interests are in performance evaluation of computer networks, and specifically in Bloom filters, network proxying for energy efficiency, and energy efficiency of P2P networks. His expected graduation date is December 2009. His homepage is at http://www.cse.usf.edu/~mjimeno.
