A "Green TCP/IP" to Reduce Electricity Consumed by Computers

Liam Irish and Kenneth J. Christensen
Computer Science and Engineering
University of South Florida
Tampa, Florida 33620
{irish, christen}@csee.usf.edu

Abstract

This new era in engineering calls for a greater understanding of the environmental impacts of modern technology. In the year 2000 it is estimated that Personal Computers in the USA alone will consume 21.9 TWh of electricity. Most of this electricity will be wasted due to PC's remaining powered-on, but idle, most of the time. PC's are often left powered-on so that network connectivity can be maintained even when the PC is not actively used. This paper describes how PC's can be powered-off and network connectivity still be maintained. The design and evaluation of a new connection sleep option in a "Green TCP/IP" is described. The long-term impacts of a Green TCP/IP can be very significant in terms of measured electricity savings.

1. Introduction

In the year 2000 it is estimated (see [3]) that without power management 21.9 TWh of electricity, or \$1.75 billion based on \$0.08 per kWh, will be consumed by Personal Computers (PC's). Slightly less than half of this consumed electricity is from the PC system units, the remainder from the monitors. Much of this energy will be wasted. Many studies have shown that most of the time that PC's are powered on they are not in use. For example, in [4] is it shown that 74% of PC's are on during the day, but used only about 12% of the time and 21% of PC's are on overnight and weekends with no use. To address this large electricity waste, the Environmental Protection Agency (EPA) developed the Energy Star program for PC's in 1992 (see [8]). To achieve an Energy Star logo a PC must consume less than 60 W during periods of inactivity. In 1995 an executive order signed by President Clinton required that all PC's purchased by government agencies be Energy Star compliant. However, there are no requirements that the Energy Star features must remain enabled following installation of a PC.

To achieve EPA Energy Star compliance (or "Green PC" operation), modern high-performance PC's (e.g., those based on Pentium-II processors in late 1997) must completely power-off their processors during periods of inactivity. The current software state of the PC can be saved to disk before power-off and then be reloaded on detection of keyboard or mouse activity. This is very similar to the operation of a laptop PC. However, a missing ingredient is in network accessibility. A Green PC entering a power-managed "sleeping" state presents two challenges:

1. Any existing network connections are dropped resulting in lost work and inconvenience.
2. It becomes inaccessible to other PC's or workstations wishing to access its resources.

The end-result is that Energy Star features are often disabled by the user or network administrator. In the case study in [5] it is reported that only 11% of the studied PC's were fully-enabled for Energy Star operation. To address these challenges, a "Green TCP/IP" is being developed at the University of South Florida. In this paper challenge (1) above is addressed.

The remainder of this paper is organized as follows. Section 2 is a brief background on TCP/IP connections and packet flows. Section 3 describes the connection sleep state implemented in Green TCP/IP. In section 4 an experimental demonstration of connection sleep is described. Finally, section 5 is a summary followed by references.

0-7803-4391-3/98/\$10.00 1998 IEEE

2. Background on TCP/IP Connections

Transmission Control Protocol / Internet Protocol (TCP/IP) is the predominant "higher-layer" protocol used in networks today including within Local Area Networks (LANs) and the Internet. Figure 1(a) shows client and server hosts attached to an Ethernet and to the Internet via routers. The key function of TCP/IP is to assure delivery of data between two hosts, typically a receiving client and a sending server. Assured delivery is accomplished by a formal connection between two hosts with acknowledgments sent for each successfully received data packet. Connection establishment is based on a three way handshake as shown in Figure 1(b) using Synchronize (SYN) packets followed by an acknowledgment (ACK) packet. Then, for each data packet sent, the receiving host returns an ACK in a separate packet or piggy-backed to a data packet. If the sending host does not get an ACK within a predetermined time-out period, the presumed lost data packet is retransmitted. This retransmission on time-out occurs for a fixed number of times and then the sending host drops the connection. At the completion of a connection, the two hosts exchange FINISH (FIN) packets in much the same way as the connection establishment SYN packet exchange (the FIN exchange is not shown in Figure 1(b)). When a connection is dropped, either by excessive time-outs or a FIN handshake, the resources dedicated to the connection are "cleaned-up". It is in this resource clean-up that data files opened by a client in a server may be lost. In addition, reconnection of a dropped connection may require time consuming logon procedures to be executed by a user.

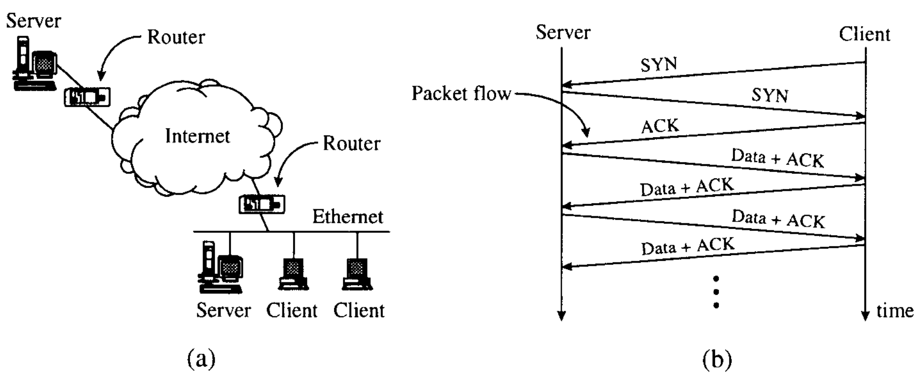

Figure 1 - (a) Network with clients and servers and (b) TCP/IP connection timing diagram

A power-managed client with existing TCP/IP connection(s) (e.g., logged-on to a server) cannot enter a sleep state without dropping its connection(s). This is because when the client is in the sleep state its system processor (where TCP/IP code executes) is stopped and thus cannot process incoming packets and transmit the needed ACK packets. One solution to this problem is to develop specialized LAN adapters that include TCP/IP processing thus allowing TCP/IP to continue operating despite a system processor power-off. However, this hardware solution is much more costly than a software-only solution and cannot be applied to already installed systems. A software-only solution is part of the Green TCP/IP under development at the University of South Florida.

3. Development of a Connection Sleep for Green TCP/IP

A key component in Green TCP/IP is a new connection sleep option. The sleep option allows a power-managed Green TCP/IP client to notify a Green TCP/IP server that it is "going to sleep". Upon this notification, the server logically keeps the connection alive, but does not send any data or ACK packets to the sleeping client. Further, the socket associated with the connection is blocked in the server to prevent excessive queueing of data to send. When the Green TCP/IP client wakes-up, it notifies the server and data flows can immediately be resumed. The changes to

TCP/IP to implement the connection sleep must be backwards compatible so that Green TCP/IP hosts can co-exist with non-Green, or regular, TCP/IP hosts. The solution to this is to create an option in the TCP header, called TCP_SLEEP, that will tell a connected server that a client is going to sleep. When the server receives a packet with this option, it will bypass all internal TCP/IP instructions which would drop the connection for that client. According to the TCP standard, a machine which does not understand an option will ignore it (see [7]), thus making the new TCP_SLEEP option backwards compatible. When the client comes back up (i.e., wakes-up due to detection of user activity), it can send a data packet on the sleeping connection and connection activity between the server and client resumes. Figure 2 shows the sending of a TCP_SLEEP packet, subsequent connection sleep, client wake-up, and resumption of normal connection activity.

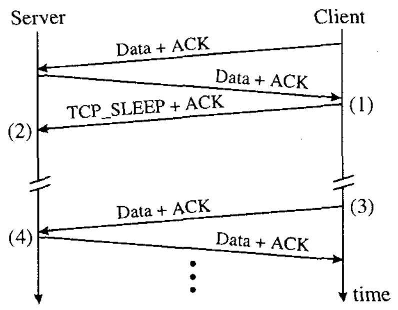

(1) = Client begins shutdown by sending TCP_SLEEP
(2) = Server blocks socket and connection enters "sleep"
(3) = Client wakes-up and sends data to server
(4) = Server receives data and resumes connection

Between (2) and (3) the client can be fully powered-off and the connection is "sleeping"

Figure 2 - Connection state diagram showing new sleep state for Green TCP/IP

An experimental implementation was produced using the Linux operating system. Linux is a complete Unix clone with source code for Intel x86 and Pentium, DEC Alpha, and SPARC based machines (see [1]). Linux is used in industry and government for experimental computer science and engineering research (see, for example, the Beowulf clustered super-computer project at NASA at [2]). For Green TCP/IP the TCP portion of the Linux kernel was modified. First, the ability of the kernel to drop connections on time-out was deleted. When a power-down was emulated (by pulling out the network cable), the server did not drop its connections. When the cable was reconnected, the machines did not start exchanging data until about two minutes had passed, approximately the length for a TCP time-out. The client was then flooded with data from the server. To fix the flooding of the client on wake-up, the server was instructed to force sleeping connections to block. A blocked connection cannot accept data from an application. When data is received from the client, the server automatically unblocks the socket. Finally, the TCP code on the clients was changed so that it sends a packet with the newly defined TCP_SLEEP option. It is expected that the methods developed with Linux can be extended to TCP/IP implementation in Windows and other operating systems and environments. However, more experimentation must be done due to the many different implementations of TCP in existence. More work must also be done in thoroughly testing the modifications to the kernel as there may be unexpected side effects visible only in certain conditions.

4. Demonstration and Evaluation of Green TCP/IP Connection Sleep

To demonstrate and evaluate the Green TCP/IP connection sleep, a small testbed consisting of two servers and one client was set-up. Figure 3 shows the testbed where one server has Green TCP/IP installed, one server has unmodified (regular) TCP/IP, and the client is also running Green TCP/IP. An experiment was conducted where two telnet sessions from client to server #1 and to server #2 were established (shown as connection #1 and #2 on the Figure 3). Telnet is an application that allows a user to interface (via a text terminal window) with a remote system. Each telnet session uses a TCP/IP connection for communication between the server and client. For each telnet session the Unix biff email notification program was executed. An incoming email to the the server would then cause a small "biff message" to flow across the network to the telnet window on the client. Once the telnet sessions were established the client was made to emulate both a power-down sequence and the sending of the TCP_SLEEP packet (see Figure 2) to the two servers. Then, the network cable to the client was removed to emulate a complete power-down of the client. During the time period of client disconnection, multiple email messages were sent to users logged-in to server #1 and server #2. When the client was re-connected to the network, emulating a power-up due to user activity, the telnet sessions were checked. In this experiment the telnet session for server #1 remained "up" and the biff email notification messages appeared correctly. However, the telnet session for server #2 (the non-green server) had "crashed" (due to TCP/IP connection time-out) and no biff messages appeared. Other applications, including remote login (rlogin) were tested and produced the same results.

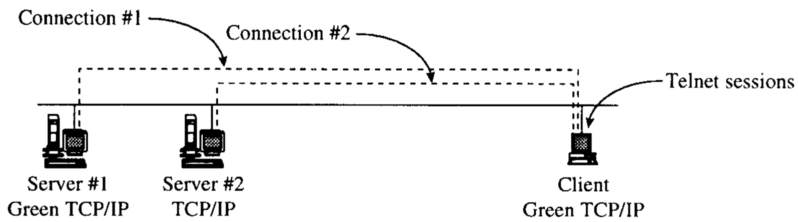

Figure 3 - Green TCP/IP test configuration

5. Conclusions

This paper has presented the development and evaluation of a Green TCP/IP that enables existing TCP/IP connections to be "put to sleep". The developed method is backwards compatible with existing non-green TCP/IP. Green TCP/IP is a first step in solving power management for network-attached PC's allowing a greater percentage of PC's to remain fully Energy Star enabled. Due to the very large numbers of PC's, even a small percentage increase in Energy Star enablement can result in very large electricity savings. Future work in connection sleep will address, 1) interfacing the sending of TCP_SLEEP to the Advanced Power Management firmware installed in all Energy Star PC's, 2) establishing a maximum sleep period to prevent "forever sleeping" connections, and 3) establishing a three-way handshake on both entering and exiting connection sleep. Funding for further work is currently being sought from the EPA.

References

[1] M. Beck, et al. LINUX Kernal Internals, Addison-Wesley, Harlow, England, 1996.
[2] Beowulf cluster computing project, NASA, 1997. URL: *http://cesdis.gsfc.nasa.gov/linux-web/beowulf/beowulf.html*.
[3] J. Koomey, M. Cramer, M. A. Piette, and J. Eto, "Efficiency Improvements in U.S. Office Equipment: Expected Policy Impacts and Uncertainties," *LBL -3738*, Lawrence Berkeley Laboratory, December 1995.
[4] G. Newsham and D. Tiller, "The Energy Consumption of Desktop Computers: Measurement and Savings Potential," *IEEE Transactions on Industry Applications*, Vol. 30, No. 4, July/August 1994.
[5] B. Nordman, M. A. Piette, and K. Kinney, "Measured Energy Savings and Performance of Power-Managed Personal Computers and Monitors," *Proceedings of the 1996 ACEEE Summer Study on Energy Efficiency in Buildings*, August 1996.
[6] W. Stevens, TCP/IP Illustrated, Volume 1, The Protocols, Addison-Wesley, Reading, Massachusetts, 1994.
[7] W. Stevens, TCP/IP Illustrated, Volume 2, The Implementation, Addison-Wesley, Reading, Massachusetts, 1995.
[8] "U.S. EPA Energy Star Office Equipment Program." URL: *http://www.epa.gov/office.html*.