An Empirical Model of Power Consumption in the NetFPGA Gigabit Router

Arun Vishwanath†, Zhi Zhao†, Vijay Sivaraman†, Craig Russell‡
†School of EE&T, University of New South Wales, Sydney, Australia
Emails: {arunv@ee.unsw.edu.au, zhi.zhao@student.unsw.edu.au, vijay@unsw.edu.au}
‡CSIRO ICT Centre, Sydney, Email: {craig.russell@csiro.au}

***Abstract*— Sustained exponential bandwidth growth of the Internet is threatened by concomitant increase in power requirements of network switches and routers. Recently, researchers have profiled the energy demand of commercial routing platforms and proposed ways to reduce network power consumption. However, the profiles are coarse-grained, and the platforms are inflexible for experimenting with new schemes for saving energy. In this paper, we consider the NetFPGA platform, which is a popular routing platform for networking research, and make three new contributions. First, we obtain fine-grained measurements of the power consumed by the NetFPGA Gigabit router over a range of packet rates and packet sizes. Second, we propose a simple model that breaks down power consumption of the NetFPGA router at the granularity of per-byte storage and per-packet processing. Lastly, we point to ways of saving network energy based on our fine-grained profile of router power consumption. We believe the results reported in this paper are a first-step towards enabling the use of fine-grained measurement data for energy-optimised network architectures.**

I. Introduction

Internet traffic has witnessed exponential growth over the past decade – it is currently in the Exabyte range (10^{18} bytes) per year, and projections show it will reach Zettabytes (10^{21} bytes) within the next 5 years [1]. To cope with such high traffic demand, a typical telecommunications rack today switches data at terabits-per-second speeds. This however comes at the cost of it drawing anywhere between 10-20 KiloWatts of power, which is five to seven times higher than ten years back. A major survey of network operators and equipment vendors was recently conducted by Ovum, and the report released in Mar 2009 states that by 2020 telcos would be switching petabits-per-second of traffic and would consume a MegaWatt of power. This trend presents grave concerns, mounting pressure on ISPs and equipment vendors to "go green". A consortium called GreenTouch [2] has recently been established, which promises to improve the energy efficiency of the communication technology sector by a factor of 1000 from current levels by 2015.

Addressing the energy efficiency challenges in wireline networks is receiving considerable attention in the literature [3], [4], [5], [6]. Some researchers have measured energy consumption of commercial routing platforms, using which they have developed models for saving energy via techniques such as turning off ports / linecards / routers, adapting link rate, and power-optimising network design. In this paper, we focus on the NetFPGA experimental Gigabit routing platform. Our reasons for this are twofold. First, commercial routers do not offer sufficient flexibility (or mechanisms) for prototyping and implementing novel power-optimised network algorithms, whereas the NetFPGA router is an open reprogrammable hardware platform that is gaining popularity amongst networking researchers world-wide, allowing new mechanisms (for power savings) to be easily prototyped and evaluated. Second, commercial platforms provide only coarse-grained measurements of power (typically obtained by measuring the power consumption of the entire device, or via command line interface commands), whereas in this paper we focus on obtaining fine-grained energy measurements for per-byte storage and per-packet processing, which typically cannot be measured with high confidence in commercial routers today.

In this paper, we break down the total power consumed by the NetFPGA Gigabit router into more fundamental components, and quantify the power consumed for each of these operations. In this context, our contributions in this paper are as follows. First, we obtain fine-grained measurements at the granularity of the energy required to store/process a packet through a Gigabit Ethernet router. To the best of our knowledge, ours is the first work that gathers such fine-grained information. Second, we develop a simple model that is able to accurately predict the power consumption of the router for different packet rates and packet sizes. Finally, we speculate on how these numbers can shed further light on designing energy efficient network architectures. We believe our results in this study complement other measurement data reported in the literature, and enables prototyping of novel power efficient network designs.

The rest of the paper is organised as follows. In Section II, we describe the NetFPGA Gigabit router used in our experiments and outline the life of a packet through the router. The experimental setup and the measurement results are presented in Section III, while a simple analytical model to estimate the power consumption is developed in Section IV. We conclude the paper in Section V with pointers to directions for future work.

II. The NetFPGA Gigabit Router

The NetFPGA is a PCI-based hardware platform developed by Stanford University, which consists of a fully programmable FPGA based core with four Gigabit Ethernet interfaces, and functions as an IP router [7]. The FPGA handles all the packet-processing tasks in *hardware* (such as switching and routing), contains 4.5 MB of SRAM (to store the routing and ARP tables), and 64 MB of DRAM

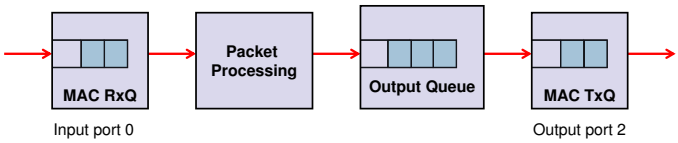

Fig. 1. Life of a packet through the NetFPGA router

(to serve as packet buffers). The experimental work in this paper uses the NetFPGA revision 2 board and the standard reference router bit file. The bit file is synthesised in such a way that packets are buffered in the SRAM, not the DRAM. The NetFPGA has been used extensively in numerous experimental studies and to tackle various research problems – study of router buffer sizing, clean-slate network architecture and designs, to name a few. For more information, refer [7].

A. Life of a Packet Through the NetFPGA Router

Fig. 1 shows a high-level view of the life of a packet inside the NetFPGA router. Let us first assume that a packet is *received* on port 0 via the Ethernet MAC receive queue (MAC RxQ). Next, the output port lookup module inside the FPGA performs the following *packet processing* steps: first, it does a longest prefix match on the routing table to decide the output port the packet is destined to (in this case port 2), second, it obtains the next-hop MAC address by performing an ARP lookup, and finally, the source/destination MAC addresses in the packet header are modified accordingly, the TTL (time-to-live) is decremented, and the checksum is updated. The packet is then buffered in the respective output queue. While all these operations are being performed, the packet is *stored* in memory (on-chip and SRAM). For transmission, it is extracted from the output queue (SRAM) and fed to the Ethernet queue MAC TxQ port 2 ready to be forwarded.

III. Experimental Setup and Results

The NetFPGA router is a card that plugs into the PCI slot in a standard Linux based PC. In our first attempt we measured the power consumption of the whole PC that houses the NetFPGA router, and found it fluctuating in the range of 70-80 Watts (without any traffic). When we injected traffic through the NetFPGA router, there was no perceptible change in power consumption; most likely it was hidden in the noise from other components in the PC such as moving disks, fans, etc. Other authors also face this problem, such as in [5] the base-line power consumption of the Cisco GSR and 7507 are upwards of 200 Watts, making it difficult to quantify the impact of packet traffic with any confidence.

We therefore isolated the power consumption of the NetFPGA router card by mounting it on an Ultraview PCI Smart Extender PCIEXT-64U riser card [9], which acts as an interface between the Linux PC and the NetFPGA router. The PCI riser card provides power leads to which we connected a Cleverscope CS328A oscilloscope [10], thus allowing us to measure the power consumed by the NetFPGA router card accurately. To generate traffic, we used the very high precision IXIA [8] hardware traffic generator with four Gigabit Ethernet ports. Constant bit rate traffic from the IXIA is fed into Ethernet port 0 of the NetFPGA, and is routed out of port 2 back to the IXIA sink.

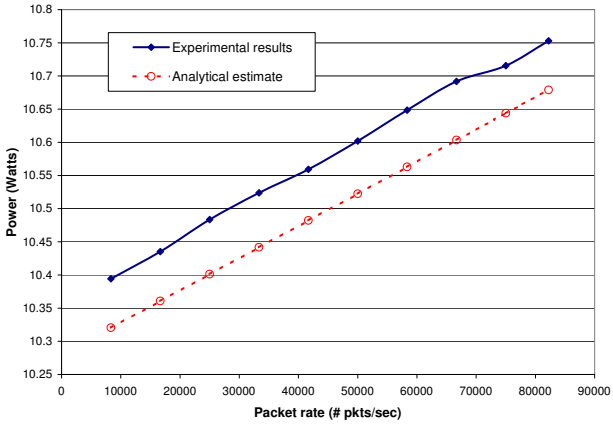

Fig. 2. Power consumption versus packet rate (1500 byte packets)

A. Impact of Varying the Packet Rate

In the first set of experiments, we vary the packet rate (equivalently, we vary the data rate from 100 Mbps to ≈ 1 Gbps) of the traffic feeding into the router while keeping the packet length constant at 1500 bytes. The solid curve in Fig. 2 shows the (measured) router power consumption as a function of the input rate (in packets/sec), and indicates that the power consumed varies linearly as the input rate increases – from 10.394 Watts to 10.753 Watts (an increase of $\approx 3.5\%$) as the data rate increases by a factor of ten.

B. Impact of Varying the Packet Size

In the next set of experiments, we vary the packet size of the ingress traffic while keeping the packet rate constant at 60000 packets/sec, and plot (using solid line) in Fig. 3 the power consumed for different packet size settings. Although we keep the packet rate constant, the offered load (in terms of Mbps) increases with packet size, and not surprisingly, the power consumed also increases fairly linearly, from 10.304 Watts to 10.572 Watts, when the packet size increases from 100 bytes to 1500 bytes.

Using the above experimental results, in the next section we develop a simple linear model to quantitatively capture the power consumed by each of the high-level stages that a packet goes through within the router.

IV. A Simple Analytical Model

Denote by P the total power consumed by the router for any given experimental setting. Then the contribution of each of the stages to the overall power consumption is as follows. First, the size of the packet determines the energy needed to transmit/receive it on the Ethernet link. Larger packets in general require higher transmission/reception energies than smaller packets. Second, the storage energy increases linearly with packet size, since larger packets occupy more memory space. Third, the packet processing energy, which is the energy consumed for header processing, ARP lookups, etc., is independent of the packet size, since these operations have to be performed on each and every packet, regardless of its size. Finally, we observe that there is a constant (base line) power consumed by the router when it is idle, and not subjected to any traffic load.

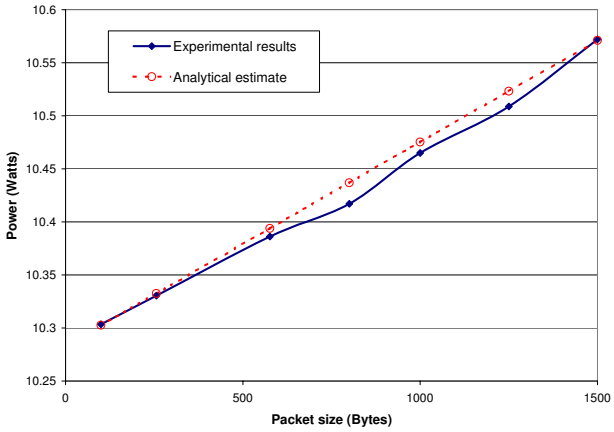

Fig. 3. Power consumption versus packet size (60000 packets/sec)

The above intuition helps us develop a simple linear model that integrates the various components as shown below.

$$P = C + NP_{proc} + NL\left(P_{rx} + P_{sto} + P_{tx}\right) \quad (1)$$

where

- C is the constant base line power,
- N is the packet rate, (i.e., the load in terms of pkts/sec),
- L is the packet size in bytes,
- P_{rx} (P_{tx}) is the Ethernet receive (transmit) energy/byte,
- P_{proc} is the energy required to process every packet, and
- P_{sto} is the storage energy (per byte).

A. Estimating the Individual Energy Components

The above equation along with the two experimental results outlined in Section III can be used to estimate the values of P_{tx}, P_{rx}, P_{sto} and P_{proc}. Using the total power consumption for the 1500 byte and 100 byte packet settings from Fig. 3, (1) can be written as

$$C + 60000P_{proc} + 60000 \cdot 1500\left(P_{rx} + P_{sto} + P_{tx}\right) = 10.572$$
$$C + 60000P_{proc} + 60000 \cdot 100\left(P_{rx} + P_{sto} + P_{tx}\right) = 10.304$$

Subtracting one from the other and simplifying results in

$$P_{rx} + P_{sto} + P_{tx} = 3.196 \quad (2)$$

The above equation says that the energy required to receive, store and transmit one byte is 3.196 nJ.

Similarly, taking the difference in power between the two end points in Fig. 2, and simplifying yields

$$P_{proc} + 1500\left(P_{rx} + P_{sto} + P_{tx}\right) = 4.852 \cdot 10^{-6} \quad (3)$$

The processing energy per packet P_{proc} can now be obtained by substituting (2) in (3) which results in $P_{proc} = 57.331$ nJ.

It has been reported in [11] that for commercial off-the-shelf Cat 5e Ethernet cables, the transmission power is approximately 62 mW. Since the router operates at 1 Gbps, a single byte is transmitted in 8 ns, therefore the energy required to transmit one byte is 0.496 nJ. If we assume the receive energy to be the same as the transmit energy per byte, then $P_{rx} = P_{tx} = 0.496$ nJ. Plugging this in (2) results in P_{sto} the energy required to store one byte to be 2.204 nJ. We can now use these values along with the experimental results to obtain an estimate for C, the constant base line router power consumption; it turns out that C is ≈ 10.28 Watts.

Figures 2 and 3 also show, via dotted lines, the total power consumption estimated using our linear model. It can be seen that the analytical and experimental results match well, suggesting that (a) the model is realistic, (b) the numbers we have obtained for P_{rx}, P_{tx}, P_{sto} and P_{proc} are precise, and (c) the model can estimate the overall power consumption of the router accurately.

B. Implications of the Energy Numbers

We now speculate on how our results can be used in the design of energy efficient networks. First, to store a typical 1500 byte IP packet, our numbers indicate that it would require about 3.306 mJ, which is about sixty times the packet processing energy of 57.331 nJ. This indicates that for large packets, the storage energy is non-negligible compared to the processing energy, at least for simple router configurations. This opens up the possibility of adapting the storage size to meet energy constraints. Second, the constant base line power C is the dominant factor in (1), however as memory and processors in the future become load-proportional, the impact of the constant will diminish while other factors such as P_{proc} and P_{sto} will begin to dominate. Our work opens the doors for understanding how router power consumption will change if the underlying components were to become load-proportional. Finally, TCP Reno/New-Reno are the most prevalent transport protocols used in the Internet today. However, these are not optimised for minimising the power consumption of data transfers. Our work allows us to measure the energy footprint of TCP file transfers and provides a handle for prototyping new energy efficient transport protocols.

V. Conclusions and Future Work

Power consumption of the Internet is becoming increasingly important, and researchers are working on various schemes to improve its efficiency. In this paper, we undertook a preliminary experimental study using the NetFPGA Gigabit router and obtained fine-grained measurements for the per-packet (per-byte) processing (storage) energy. We also proposed a simple model for the power consumption and showed that it matches well with the experimental results. We are currently working on quantifying the energy impact of packet buffers in the presence of congestion.

References

[1] Cisco Whitepaper: "Approaching the Zettabyte Era," Jun 2008,
[2] Green Touch – http://www.greentouch.org
[3] M. Gupta and S. Singh, "Greening the Internet," *Proc. ACM SIGCOMM*, Germany, Aug 2003.
[4] S. Nedevschi et al., "Reducing Network Energy Consumption via Rate-Adaptation and Sleeping," *Proc. NSDI*, USA, Apr 2008.
[5] J. Chabarek et al., "Power Awareness in Network Design and Routing," *Proc. IEEE INFOCOM*, USA, Mar 2008.
[6] P. Mahadevan et al., "A Power Benchmarking Framework for Network Devices," *Proc. IFIP NETWORKING*, Germany, May 2009.
[7] NetFPGA Gigabit Router – www.netfpga.org
[8] Ixia traffic generator – http://www.ixiacom.com
[9] Ultraview PCI Smart Extenders – http://ultraviewcorp.com/displayproduct.php?part_id=4&sub_id=2
[10] Cleverscope oscilloscope – http://www.cleverscope.com/products/cs328a.php
[11] Y. Chen and R. H. Katz, "Energy Efficient Ethernet Encodings," *Technical Report No. UCB/EECS-2009-68*, University of California at Berkeley, May 2009.