Available online at www.sciencedirect.com
ScienceDirect
ELSEVIER
June 2012, 19(Suppl. 1): 94–99
www.sciencedirect.com/science/journal/10058885
The Journal of China Universities of Posts and Telecommunications
http://jcupt.xsw.bupt.cn

Power consumption analysis of a NetFPGA based router

Feng Guo, Olga Ormond, Leonardo Fialho, Martin Collier, Xiaojun Wang (✉)

The Rince Institute, School of Electronic Engineering, Dublin City University, Dublin 09, Ireland

Abstract

For both economic and environmental reasons, energy efficiency is becoming increasingly important in the design of next generation networks (NGN). The energy efficiency improvements for network components can mainly be achieved by the support of smart standby and/or frequency scaling. This paper describes fine-grained power measurements of the peripheral component interconnect (PCI)-based network field-programmable gate array 1 gigabit (NetFPGA 1G) reference router when scaling the frequency of router core logic and static random access memories (SRAMs) between 125 MHz and 62.5 MHz. This paper presents the power consumption of a NetFPGA 1G reference router under different scenarios. Results show that by reducing the frequency from 125 MHz to 62.5 MHz, under a user datagram protocol (UDP) traffic load of 400 Mbit/s, 12.23% of power can be saved with the same quality of service (QoS), i.e. no packet loss in either case. Moreover, aggregating the traffic and rerouting the packets can save relatively high amount of energy. For example, our results show that 19.77% of power consumption can be saved by aggregating four 100 Mbit/s links into two 200 Mbit/s links.

Keywords NetFPGA 1G, energy efficiency, frequency scaling

1 Introduction

With the development of NGN, it is widely believed that, energy efficiency should be taken into consideration in wired network infrastructures. Ethernet dominates wired communications technology for local area networks (LANs), with over 3 billion interfaces installed worldwide [1]. Early works such as Ref. [2] estimates that for European Internet service providers (ISPs) the overall network power consumption in 2010 was approximately 21.4 TWh. This work also predicts a huge increase of about 35.8 TWh in 2020 if no green technologies are used. Two main driving forces that motivate the requirement of energy efficient technologies are: the protection of the environment, meeting CO_2 emissions reduction targets; and economics, fulfilling the reduction of maintenance cost for both operators and users, while ensuring end-to-end QoS.

Currently, switches and routers do not include comprehensive energy consumption values. Most device specification sheets only report the maximum rated power. This information is insufficient to understand the actual energy consumption of a network device. As shown in the rest of this paper, the actual energy consumed by wired routers depends on various factors, such as: 1) working speed, 2) number of active Ethernet ports and 3) traffic load. Moreover, most network links and devices are provisioned for busy or rush hour load, which typically exceeds their average utilization. It means that for a large percentage of the operation time these links are over-provisioned.

In order to have more insight into the power consumption of wired network devices, especially for future design, we examine a NetFPGA 1G board [3]. This board provides a quick and easy way to implement a custom router equipped with four 1 Gbit/s Ethernet ports. It is an open source and modular platform, which permits developers to share and build on each other's projects and intellectual property building blocks. Based on this hardware platform, we: 1) measure the power consumption under varying traffic loads, 2) present results for

Received date: 12-05-2012
Corresponding author: Xiaojun Wang, E-mail: xiaojun.wang@dcu.ie
DOI: 10.1016/S1005-8885(11)60462-7

performance degradation when operating at the scaled down frequency, and 3) use this to examine where more energy savings would be possible in future designs.

From our benchmarking power measurements, the results show that:

1) At the higher frequency (125 MHz) the NetFPGA 1G router consumes more energy than for 62.5 MHz operation, for similar traffic loads to the limit of 2 Gbit/s.

2) The NetFPGA 1G router power consumption increases linearly with the number of active ports (from 0 to 4).

3) The NetFPGA 1G router power consumption is proportional to its traffic load.

The rest of this paper is organized as follows. Sect. 2 describes the NetFPGA 1G board and the reference router. Sect. 3 illustrates the experiment setup for measuring the power consumption data. In Sect. 4, we present the measurements and analyze the power consumption of the NetFPGA 1G router. Sect. 5 concludes the paper and points to possible directions for future work.

2 NetFPGA 1G board and the reference router

The NetFPGA 1G board is an open source and low cost reconfigurable hardware platform optimized for use as a high-speed networking router. The board itself is a PCI card that consists of a small Xilinx Spartan II FPGA for the control logic from the PCI interface to the host machine and a large Xilinx Virtex-II Pro FPGA for user defined logic programming. The Spartan II FPGA works at a fixed frequency. The Virtex-II Pro clock, which is the core logic clock, can be toggled between 125 MHz and 62.5 MHz [4]. Two SRAMs run synchronously with the core logic clock at either 125 MHz or 62.5 MHz.

The NetFPGA 1G board can be set up as a reference router. According to the reference router logic, incoming packets go through 4 steps into 5 stages in the packet processing pipeline, as shown in Fig. 1: 1) from the receive queues at stage 1 to the input arbiter at stage 2, 2) from the input arbiter to the output port lookup at stage 3, 3) from the output port lookup to the output queues at stage 4, and finally 4) from the output queues to the transmit queues at stage 5.

The NetFPGA 1G board provides four external network interfaces to the host machine operating system: nf2c0, nf2c1, nf2c2 and nf2c3. The board itself is connected to the host machine motherboard through a PCI slot. The NetFPGA interfaces parameters such as Internet protocol (IP) and media access control (MAC) addresses could be configured using user-level software called SCONE [5]. This software communicates with the NetFPGA 1G board using its corresponding kernel driver. To use the NetFPGA 1G board as a reference router, the host machine operating system should not be used to configure the board's interface. Otherwise packets from the external interfaces would be hijacked by the host machine operating system. These hijacked packets would pass through the operating system kernel IP stack and update its routing and address resolution protocol (ARP) table ignoring the board routing functionality. For that reason, packets from the host machine should be routed to the NetFPGA interfaces through another network interface existing in the host machine. Thus, the stand alone host machine equipped with a NetFPGA 1G board is only set up to be used as a packet forwarder.

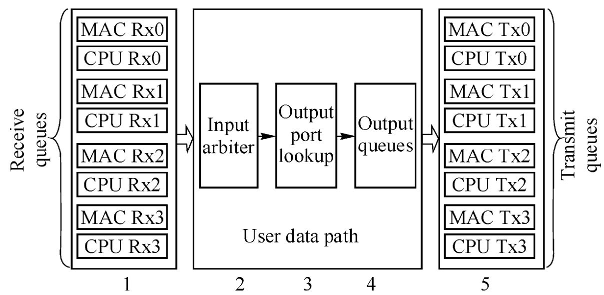

Fig. 1 NetFPGA 1G reference router pipeline

3 Experiment setup

We have set up the environment depicted in Fig. 2. This environment is composed of: 1) a power measurement device, 2) four client and server PC hosts and 3) a host machine equipped with the NetFPGA 1G board.

The power measurement computer uses a LabJack U6 [6] device to automatically collect power consumption data from the NetFPGA 1G board. Measurements are made through an Ultraview PCI bus extender [7], as shown by point 2 in Fig. 3. The LabJack U6 is an universal serial bus (USB)-based measurement and automation device that provides readings from the Ultraview PCI bus extender. A powerful measurement software running on the power measurement computer, called Labview, automatically collects and loads the power consumption data from LabJack U6. This is the accurate and fine-grained measurement method [8] we have used in this paper. Another option is to take the measurements of the global system (including the entire personal computer (PC) and the NetFPGA 1G board) through the main power supply cable (point 1 in Fig. 3) with a device like an owl energy monitor (OWL) [9] wireless electricity monitor and then estimating the power consumption by measurement comparison of the system without the NetFPGA 1G board running. This second method should not be used, since it does not provide the same accuracy as the first.

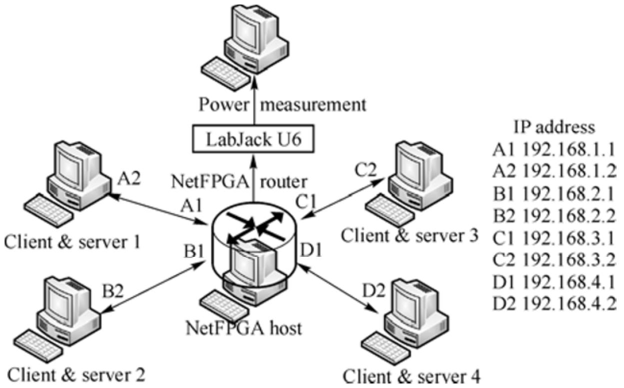

Fig. 2 Experiment topology

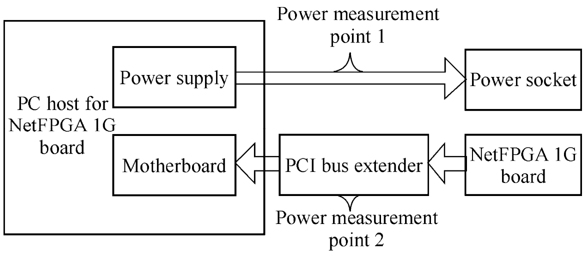

Fig. 3 Two power measurement methods that leads to two different power measurement points

The four PCs act as traffic sources and sinks, and are connected one to each of the four NetFPGA 1 Gbit/s ports. We have used Iperf [10], which is a testing tool for network performance measurement. Iperf have been used on experiments to generate varying UDP traffic among the four PCs.

The host machine PC is used to host the NetFPGA 1G reference router. In order to provide accurate power consumption data, there are no other applications that use the NetFPGA 1G board running on this PC. The host PC hardware is a Dell Optiplex 780 equipped with an Intel Core2® Duo processor model E7500 running at 2.93 GHz. It has 4 GB of Samsung DDR3 @ 1 333 MHz memory. The operating system running is CentOS version 5.5.

To set up the packet routing environment we have used the SCONE software to perform IPv4 forwarding, handle ARPs and various Internet control message protocol (ICMP) messages. SCONE has telnet (port 23) and hyper text transfer protocol (HTTP) (port 8080) services to handle router control. It also implements a subset of Pee Wee open shortest path first (OSPF) (PW OSPF). SCONE configures the NetFPGA 1G board with the MAC and IP addresses of the four interfaces and routing and ARP tables onto the NetFPGA 1G board, which hardware accelerates the forwarding path.

4 Measurements

Some initial power measurements were taken with an OWL energy meter to estimate the power consumption of the NetFPGA 1G router. These measurements show that the power consumption of the host PC for the NetFPGA 1G router fluctuates among four fixed values (32 W, 48 W, 64 W and 128 W). These fluctuations most likely come from the host PC components such as digital versatile disk (DVD) drives and fans. To isolate the power consumption of the NetFPGA 1G router, we use the power measurement environment described in Sect. 3.

To find out exactly how much power is consumed by a NetFPGA 1G router running at 125 MHz and at 62.5 MHz, we ran three groups of tests: 1) minimum power when the NetFPGA 1G board is powered on but before the reference router executable is loaded; 2) baseline power when the NetFPGA 1G board works as a router in an idle state with no ports activated; 3) power for varying traffic loads when all the NetFPGA Ethernet ports are active and working under varying traffic load. For each test run below, we have measured 10 groups of samples. For each group, tests were 10 min in duration and we took 10 samples per second. For simplicity, we only show one group of samples in each figure because we have 60 000 samples in total for each measurement. In each figure, the upper ‘x’ curve presents the information when NetFPGA 1G reference router works at 125 MHz and the lower ‘+’ curve for 62.5 MHz operation.

4.1 Minimum power

The minimum power is the power consumed when the NetFPGA 1G board is powered on with no ports active and no traffic load. If the NetFPGA 1G board is to function as a normal router, it needs to first download the reference router bitfile into the core Virtex-II Pro FPGA first. In this case, we plugged the NetFPGA 1G board into the PCI bus extender to power on the NetFPGA 1G board but without downloading the reference router bitfile. As shown in

Fig. 4, the upper 'x' curve shows the minimum power consumption of the NetFPGA 1G reference router when the core Virtex-II Pro FPGA works at 125 MHz and the lower '+' curve for 62.5 MHz. On average, the minimum power of the NetFPGA 1G reference router working at 125 MHz was 4.162 7 W, and 4.109 9 W for 62.5 MHz. The difference in minimum power consumption between 125 MHz and 62.5 MHz was only 0.052 8 W. This difference is tiny and can be explained by the fact that most components of the NetFPGA 1G board are not operational at the time of measurement, especially the core logic Xilinx Virtex-II Pro FPGA and two SRAMs.

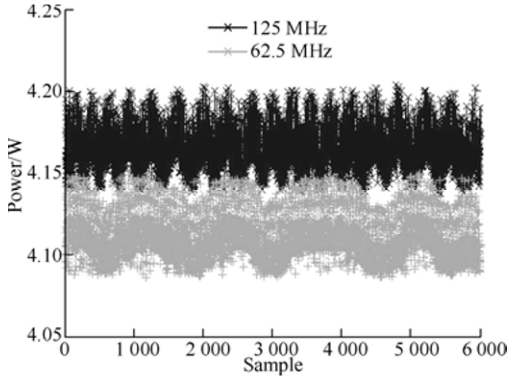

Fig. 4 Minimum power consumption for 1 group of samples

4.2 Baseline power

The baseline power is the power consumed by the NetFPGA 1G board with the reference router bitfile downloaded into its core Virtex-II Pro FPGA to work as a normal router in idle state. Idle here means that none of the Ethernet ports are activated and no traffic is involved. As shown in Fig. 5, the average baseline power of the NetFPGA 1G router working at 125 MHz was 6.572 7 W, while it is 4.661 8 W for 62.5 MHz. This 1.910 9 W difference can mainly be attributed to the difference in power consumption for the core Virtex-II Pro FPGA and two SRAMs operating at 125 MHz and at 62.5 MHz.

Following the baseline power measurements, we plugged the Ethernet cables into each NetFPGA 1G Ethernet port to activate the ports. We did this by adding one cable at a time and no traffic was involved. Results show that the power consumption increases linearly with the number of active ports on the NetPFGA 1G router. Activating each port leads to an additional~1W in power consumption, as shown in Table 1. This linear increase in power consumption can mainly be attributed to the operations of physical layer (PHY) and MAC components for each activated connection, such as enabling the link and implementing the carrier sensing.

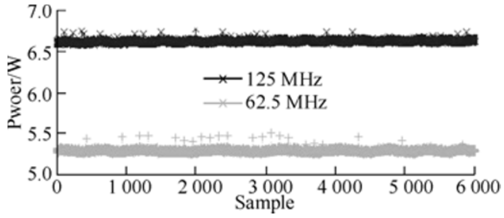

Fig. 5 Baseline power consumption for 1 group of samples

Table 1 Power consumption of varying active ports with no traffic

Number of active ports	Power consumption /W	
	125 MHz	62.5 MHz
0	6.572 7	4.661 8
1	7.610 5	5.648 8
2	8.676 8	6.652 3
3	9.705 6	7.671 5
4	10.703 6	8.688 1

4.3 Power for varying traffic loads

This test aims to measure the power consumption of the NetFPGA 1G board while all the four interfaces are active and working at varying traffic loads. The packet generator software is used to generate different traffic loads with the default packet size of 1 470 B, in the order of 100 Mbit/s, 200 Mbit/s, 300 Mbit/s, 400 Mbit/s, 500 Mbit/s, 600 Mbit/s, 700 Mbit/s, 800 Mbit/s, 900 Mbit/s and 1 Gbit/s. In this scenario, the equipment was setup as shown in Fig. 2, with the NetFPGA host PC working as a forwarding router and four PCs simultaneously sending and receiving UDP traffic. Each client/server PC uses one 1 Gbit/s Ethernet port of the NetFPGA router.

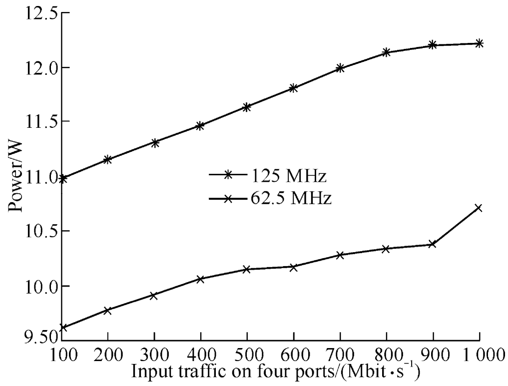

Fig. 6 Power consumption for varying input loads at both frequencies

Varying input traffic for each port was generated by Iperf software and Iperf was also used to collect the received bandwidth. In addition, we also got packet loss data from ping statistics. It is clear in figure 6 that there is more power consumed when the NetFPGA works at a higher frequency and also an increase in power consumption with increase in traffic load.

Currently the NetFPGA 1G board provides maximum 4 ports at 1 Gbit/s when the core logic is 125 MHz. If we scale down the frequency by 50% from 125 MHz to 62.5 MHz, the bandwidth performance for each port should also degrade to approximately half. That means that the NetFPGA 1G can only provide a maximum of 500 Mbit/s for each of the 4 ports when the core logic is 62.5 MHz. As seen in Fig. 7, when operating at 62.5 MHz the actual output bandwidth reaches a maximum of 480 Mbit/s even for input traffic loads above 500 Mbit/s.

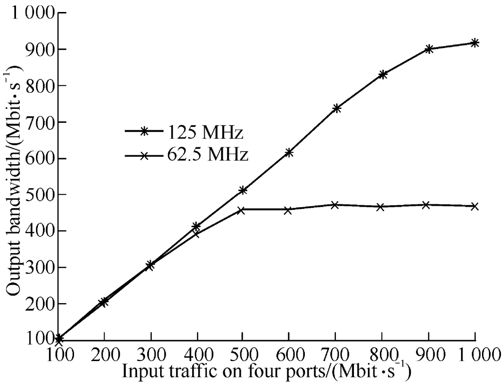

Fig. 7 Output bandwidth for varying input loads at both frequencies

As shown in Fig. 8, when the NetFPGA 1G router is operating at the lower frequency (62.5 MHz), the router can't handle traffic load greater than 500 Mbit/s and there will be very visible degradation in the order of 5.75% packet loss for 500 Mbit/s, 32% for 600 Mbit/s, 40% for 700 Mbit/s, 54.5% for 800 Mbit/s, 56.75% for 900 Mbit/s and 59.5% for 1 Gbit/s, which is unacceptable.

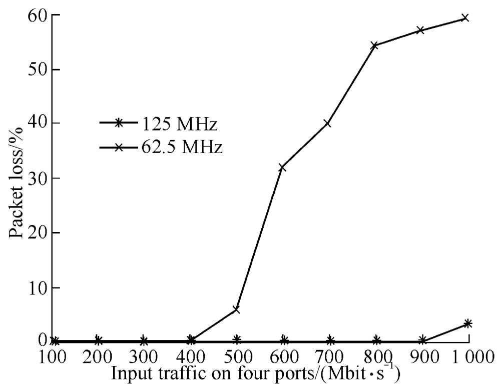

Fig. 8 Packet loss for varying input loads at both frequencies

In comparison, the packet loss for the 125 MHz is acceptable for most traffic and it only shows slight packet loss as the connections reach their capacity (1 Gbit/s). Table 2 summarizes our results for the power consumption and packet loss for varying traffic loads at both frequencies.

Table 2 Power consumption and packet loss of NetFPGA 1G reference router

Input traffic	Power consumption/W		Packet loss (%)	
(4 portsactive)/ (Mbit·s^{-1})	125 MHz	62.5 MHz	125 MHz	62.5 MHz
100	10.970 3	9.622 2	0	0
200	11.141 0	9.769 9	0	0
300	11.298 6	9.909 8	0	0
400	11.452 8	10.051 6	0	0
500	11.623 5	10.139 2	0	5.75
600	11.792 9	10.174 2	0	32.00
700	11.967 8	10.277 1	0	40.00
800	12.122 8	10.332 5	0	54.50
900	12.188 1	10.366 8	0	56.75
1 000	12.206 6	10.713 2	3	59.50

5 Conclusions

From our results, it can be seen that scaling down the frequency from 125 MHz to 62.5 MHz for a traffic load of 400 Mbit/s, can save up to 12.23% of power consumption while still ensuring the same guaranteed end-to-end QoS. Moreover, more power can be saved by turning off network ports [11] and by rerouting traffic to other ports when the traffic is low. Each port disabled can save ~1 W at each router (i.e. turning off the ports at both ends of the connection will result in an overall saving of ~2 W). Turning off a router when all ports are off can save 4 watts while scaling down the frequency to half while keeping all ports active can save a maximum of only 1.823 W. Thus, for the NetFPGA 1G reference router implementation at least, rerouting and disabling Ethernet ports is a much more effective way to reduce power consumption, and in comparison (of the power consumption amount) frequency scaling is only an add-on.

Since the NetFPGA is a research prototype, future work will be linked to commonly use commercial routers. By comparison with the power consumption of the host PC, the NetFPGA 1G router consumes a relatively small amount of power. We will make further efforts to modify and evaluate FPGA-based network interface cards to support dynamic frequency scaling and smart standby mode and examine the best energy saving router policies to adapt as part of the ECONET project.

Acknowledgements

This work was supported part by the European FP7 ECONET

Project (258454). The research is co-funded by the China Scholarship Council. The authors would like to thank the contributions of Xiangjun Shi for collecting power measurement data. Thanks also to our Italian collaborators in LightComm for useful discussion on the power measurements.

References

1. Christensen K, Reviriego P, Nordman B, et al. IEEE 802.3 az: the road to energy efficient Eethernet. IEEE Communications Magazine. 2010, 48(11): 50−56
2. Bolla R, Bruschi R, Christensen K, et al. The potential impact of green technologies in next-generation wireline networks: is there room for energy saving optimization? IEEE Communications Magazine. 2011, 80−86
3. Lockwood J W, McKeown N, Watson G, et al. NetFPGA-an openplatform for gigabit-rate network switching and routing. IEEE International Conference on Microelectronic Systems Education, 2007 (MSE’07). 2007: 160−161
4. Gibb G, Lockwood J W, Naous J, et al. NetFPGA-an open platform for teaching how to build gigabit-rate network switches and routers. IEEE Transactions on Education. 2008, 51(3): 364−369
5. Stanford University. NetFPGA SCONE Homepage. [2012-05-01]. http://netfpga.org/foswiki/bin/view/NetFPGA/OneGig/PacketGenerator
6. LabJack Corporation. LabJack U6. [2012-05-01]. http://labjack.com/u6
7. Ultraview PCI bus extender. [2012-05-01]. http://ultraviewcorp.com/displayproduct.php?part_id=4&sub_id=2
8. Sivaraman V, Vishwanath A, Zhao Z, et al. Profiling per-packet and per-byte energy consumption in the NetFPGA gigabit router. 2011 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS). 2011: 331−336
9. 2 Save Energy Ltd. OWL+USB wireless electricity monitor. [2012-05-01]. http://www.theowl.com/index.php?page=owl-usb
10. Distributed Applications Support Team (DAST) at the National Laboratory for Applied Network Research. Iperf. [2009-10-05]. http://dast.nlanr.net/Projects/Iperf/
11. Gupta M, Singh S. Dynamic Ethernet link shutdown for energy conservation on Ethernet links. IEEE International Conference on Communications, 2007 (ICC'07). 2007: 6156−6161