Towards single Watt and nJoule/bit routing

Amos NUNGU[1], Robert OLSSON[2]. Jiannan GUO[2], Björn PEHRSON[2]

[1]DIT, P.O. Box 2958, Dar es Salaam, Tanzania

Tel: + 265 1 754535, Fax: + 265 1 754535, Email: amosnungu@dit.ac.tz

[2]KTH, Electrum 229, 164 40 Kista, Sweden

Tel: + 46 8 790 6000, Email: {roolss|jiannang|bpehrson}@kth.se

Abstract

We report on recent progress in our research on how to establish sustainable broadband markets in under-served areas, in this paper focusing on facilitating sustainable power supply by designing low power-consuming routers. We disclose the details of a new generation low power routers, providing wire-speed forwarding between Fast Ethernet connections (100Mbps) over optical fiber at below 5W. Three of these routers have been deployed in the Serengeti Broadband Network in northern Tanzania, which is now an extension of TERNET, the Tanzania NREN. We expect this development to lead to even more energy-efficient routers requiring less than 1W and 1nJ/bit forwarding.

Keywords: Low power, high performance routing, rural ICT development, broadband network

1. Introduction

We report on recent progress in our research on how to establish sustainable broadband markets in under-served areas, in this case represented by the Serengeti Broadband Network in Northern Tanzania (Nungu, A., 2011). The development towards ubiquitous network access requires innovative solutions to get remote areas included. The challenges are many, including ownership and leadership, sustainable business models, robustness and affordability of network components and poor or non-existent supply chains, including power supply.

The challenges in focus here include those associated with the establishment of sustainable power supply. These challenges include non-existent or poor power grids with harmful transients and/or frequent outages and extreme environment parameters, such as temperature with ranges in which traditional batteries do not perform well and have a short life expectancy.

Our approach is to minimize the power (W) required to meet the demand for capacity (bps) and to use innovative power source and power storage technologies. We also strive at minimizing the energy/capacity ratio (J/bit or Ws/bit).

Our method is iterative systems design. To demonstrate feasibility and provide a benchmark for comparison of commercial products, we use free open source software and standard commercial off-the-shelf hardware components.

The main results reported are associated with the backbone router power requirements. We have previously set the forwarding capacity requirement to 1Gbps wireline speed over optical fiber and have arrived at a power requirement of 20-25w (Nungu, A. et al, (2012a, 2012b).

By lowering the capacity requirement to 100Mbps, we have now reduced the power requirements to less than 5W, while maintaining the power/capacity ratio roughly at the same level (<50 nJ/bit).

The current bottleneck in available hardware components is the I/O-bus. Since all available components for design of systems requiring less than 20W are currently all using USB2, the useful limit is 100 Mbps. Local network performance shows an impressive communication capacity of 5 Gbps using ipref3 over the loopback interface.

2. System Design

2.1 Motherboard and network interface cards

Our low power router designs include a motherboard, network interface card(s) (NIC) supporting one or more optical Fast or Gigabit Ethernet SFP ports and an integrated power supply unit, including a charge controller and provisioning of the required voltage(s) and current control.

Although our focus is on optical fibre, we have tested network interface cards for different media, including optical fiber or UTP-cable (100/1000 Mbps Ethernet SFP or RJ45 ports) and wireless (wife, VHF/UHF).

The motherboard used in the 20W version (Nungu, A. et al, 2012a, 2012b) of our low-power router is the Intel/ATOM-based Supermicro X7SPA (Supermicro X7SPA) supporting the PCIe I/O-bus, while the NIC is the Interface master Niagara 82048 (Interface masters Niagara 82048) providing four 1 GE SFP ports with digital optical monitoring and using the PCIe as the I/O bus. The idling power consumption of the Motherboard is 14W and for the NIC 8W. The differences between idling and full-load is negligible.

Figure 1: The first generation low-power router (20-25W) to the right and the second generation (4-5W) to the left. The new router has a much reduced footprint, weight and Bill of Material (BOM)

In the new low-power router, we use the fibergecko100 card (Lyconsys Fibergecko) with one 100/1000 Mbps SFP port each and USB2 as the I/O-bus. The power consumption is ~5W and the difference between idling and full-load is negligible. Using USB2 as the I/O bus is actually the bottleneck currently forcing the reduction of capacity. The USB standard (http://usb.org) was originally designed to replace a wide spectrum of serial and parallel

 ISSN 2223-7062 Proceedings and report of the 7th UbuntuNet Alliance annual conference, 2014, pp 325-332

standards for the connection of computer peripherals and power chargers for portable devices. The throughput capacity has grown from 12Mbps in USB1 (1996) to 280Mbps in USB2 (2000) to 10Gbps in USB3.1 (2013). We have currently not been able to find motherboards and network interface card supporting USB3. If and when components with the same functionality will support USB3 as I/O-bus, it will be possible to come back at wire speed forwarding at 1Gbps.

In SuperSpeed (SS) mode the data transfer can reach 5 Gbit/s which is tenfold improvement over USB 2.0 with 480 Mbit/s. Even an USB 3.1 revision called SuperSpeed+ has recently been announced with transfer rates up to 10 Gbit/s. This development promises low-power routers with very high bandwidth.

Our evaluation based on available motherboards for a design requiring less than 20W included

- Alix (Alix) based on the AMD Geode chipset, using ISA-style I/O bus, three 10/100 Mbps Ethernet RH45 ports, two USB2
- Raspberry Pi (Raspberry Pi)
- BeagleBone Black (BeagleBone Black)and
- Odroid U3 (Odroid U3/XU3a low cost and high performance development platform based on a Samsung Exynos 4412 ARM Cortex-A9 Quad Core 1.7GHz CPU with 2GByte RAM, a 10/100Mbps Ethernet with RJ-45 LAN Jack, 3 High speed USB2.0 host ports and GPIO/UART/I2C ports. The idling power of the motherboard is slightly more than 1W. The physical size of the motherboard is : 83 x 48 mm, Weight : 48g including heat sink

2.2 Power supply

The power supply unit includes solar cells and 220AC/13VDC power adapter as sources connected to a charge regulator card controlling the charging of a 12 V lead acid battery and monitoring the current levels. In the 20W version, the required voltage levels are provided via a picoPSU while in the new version a step-down converter to 5VDC is on the charge controller to power the Odroid Motherboard and Fibergecko NICs.

Due to earlier problems with the performance and lifetime of the lead-acid batteries due to high operating temperatures up to 40C, we are exploring different way of cooling the batteries. A rule of thumb says that each 8°C (15°F) rise in temperature halves the lifetime of a sealed lead acid battery. Also, the charging voltage must be adjusted to battery voltage, the higher voltage the lower charging voltage. This implies that the charge controller has to monitor battery temperature via a sensor on the battery. Many charge controllers assumes a nominal voltage, often 25C. Battery type (wet/sealed) and number of low-voltage disconnect (LVD) and the number of micro-cycles and depth of micro-cycles (typically between 20-70% of the nominal capacity) also impacts battery lifetime.

An extra challenge with solar driven systems is that the ambient temperature is higher during the sunny hours, which is when the batteries can be charged when using a solar panel. A sensor deployed at Nata site shows that ambient temperature varies between 15C and 40C, see Figure The battery life-time goes up if the temperature is lower and stable. We are now experimenting with placing the batteries underground. According to Jager, Tjapko (1982), the temperature is stable at 25C 0.5m down.

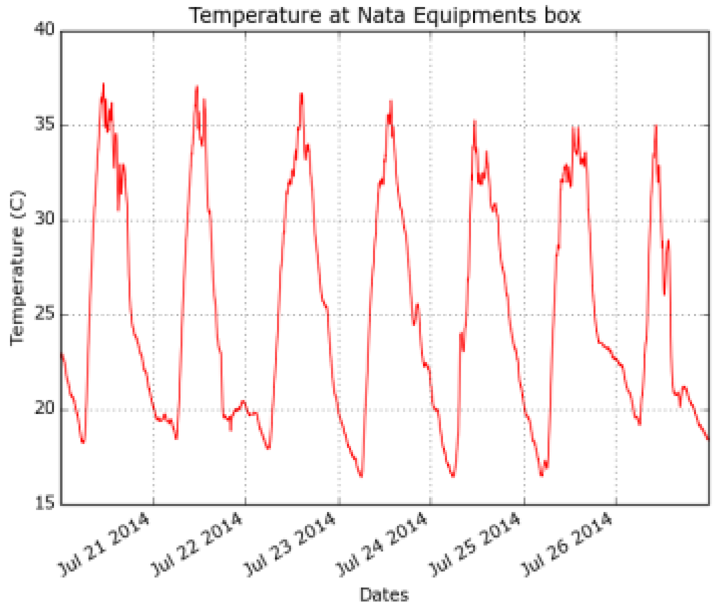

Figure 1: Temperature in Nata equipment box

As the power requirement approaches 1W, replacing traditional chemical batteries with ultra-capacitor batteries becomes economically feasible. Such batteries are less sensitive to high temperatures and can take a much higher number of discharge cycles.

2.3 Software

The operating system used in the 20W version is the Bifrost Linux-distribution. In the new version, we use a minimal version of the Debian Linux distribution, a stripped custom Linux kernel and Quagga as routing software. Some Bifrost utilities and iperf3 for performance measurement software (iperf, n.d.) have also been added. This distribution is available on request.

3. Measurements and tools

Since packet forwarding is essentially based on table lookup, memory latency and bandwidth has a high impact on network and router performance, the lower latency and higher bandwidth, the better. The memory is normally partitioned into two cache levels, L1 and L2, and main memory. Their relative sizes matter. The memory latency and bandwidth was measured using the LMBench tool set (LMBench Tools for Performance Evaluation) The results of these measurements are presented at the top in figure 2 below illustrating the power vs latency/memory bandwidth perspective

The forwarding capacity (bandwidth) was measured in a network configuration with three routers in a row with the system under test in the middle using iperf/iperf3 (iperf, n.d.) These measurements are demonstrated in the middle diagram of Figure 2.

For power measurements, a Wattson instrument indicating both voltages, current, power and energy was used.

 ISSN 2223-7062 Proceedings and report of the 7th UbuntuNet Alliance annual conference, 2014, pp 325-332

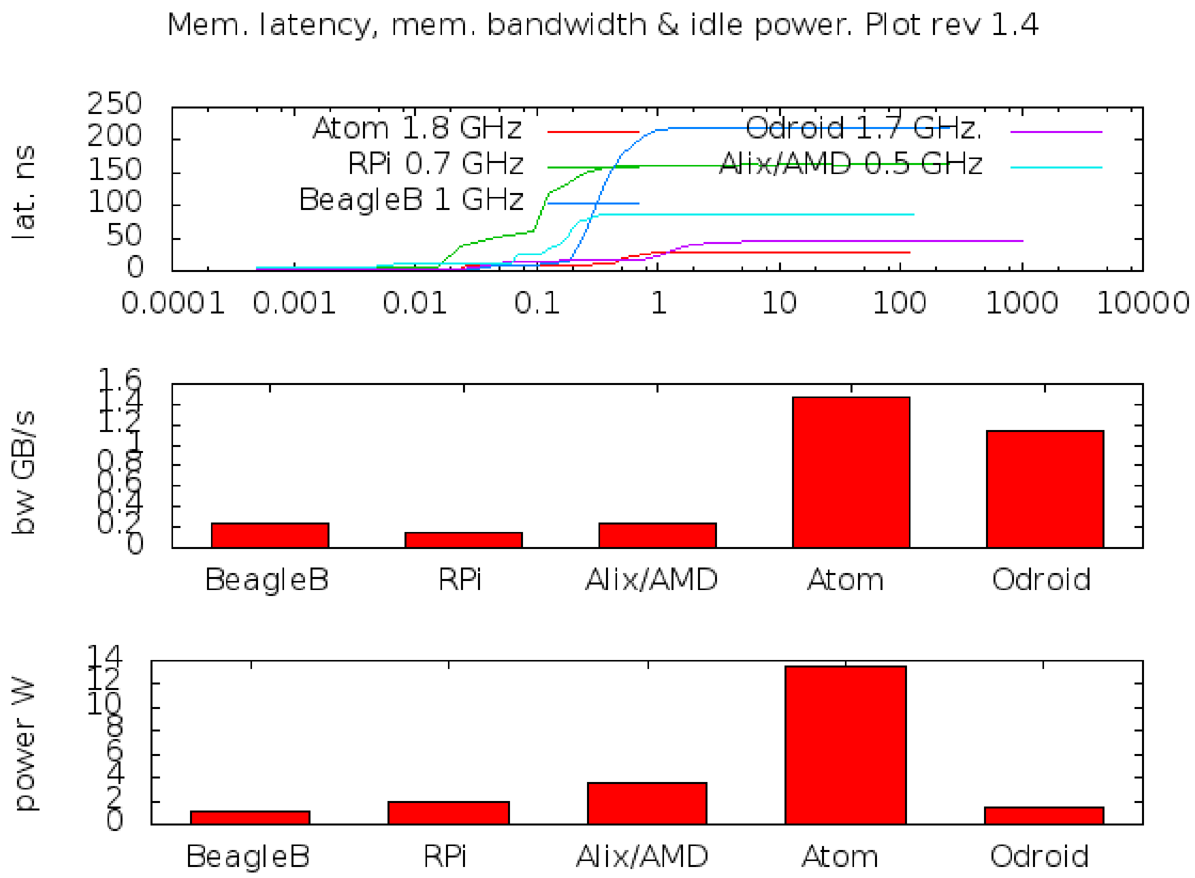

Figure 3: A benchmark of five low power platforms

The latency plot (on top), shows that the ATOM system has the lowest memory latency. The Odroid U3 is very close. We can even see Odroid U3 stays a little longer in the L2 cache.

Regarding memory bandwidth as measured with iperf (middle plot), ATOM also comes best out. Odroid U3 is, however, relatively close.

The idle power in the bottom plot is measured without NICs. The SuperMicro/ATOM-based system uses ~14W, the Alix board ~3.5W, and all three SOC-based systems ~1.4W, i.e. 10% of the ATOM power. RPi takes a little more than Odroid and BeagleBone black due to its relatively simple DC-design.

The comparison between CPU power consumption, latency/bandwidth and performance is very clear. The ATOM-based system has the best performance measure by memory latency and bandwidth, but the Odroid U3-based system is very close at 10% of the power used and at about the same power/capacity-level. It is also smaller in size and weight and more affordable.

When adding the two USB-SFP NICs (Lyconsys Fibergecko100) the total power consumed by our Odroid-based router becomes 4.57W when idling and 5.25W when fully loaded by forwarding at wire-speed ~95Mbps. The average power consumption depends on the traffic load but is normally closer to the idling power. The latency plot (on top), shows that the ATOM system has the lowest memory latency. The Odroid U3 is very close. We can even see Odroid U3 stays a little longer in the L2 cache.

Regarding memory bandwidth as measured with iperf (middle plot), ATOM also comes best out. Odroid U3 is, however, relatively close.

The idle power in the bottom plot is measured without NICs. The SuperMicro/ATOM-based system uses ~14W, the Alix board ~3.5W, and all three SOC-based systems ~1.4W, i.e. 10% of the ATOM power. RPi takes a little more than Odroid and BeagleBone black due to its relatively simple DC-design.

The comparison between CPU power consumption, latency/bandwidth and performance is very clear. The ATOM-based system has the best performance measure by memory latency and bandwidth, but the Odroid U3-based system is very close at 10% of the power used and at about the same power/capacity-level. It is also smaller in size and weight and more affordable.

When adding the two USB-SFP NICs (Lyconsys Fibergecko100) the total power consumed by our Odroid-based router becomes 4.57W when idling and 5.25W when fully loaded by forwarding at wire-speed ~95Mbps. The average power consumption depends on the traffic load but is normally closer to the idling power.

4. 4. Test and Deployment

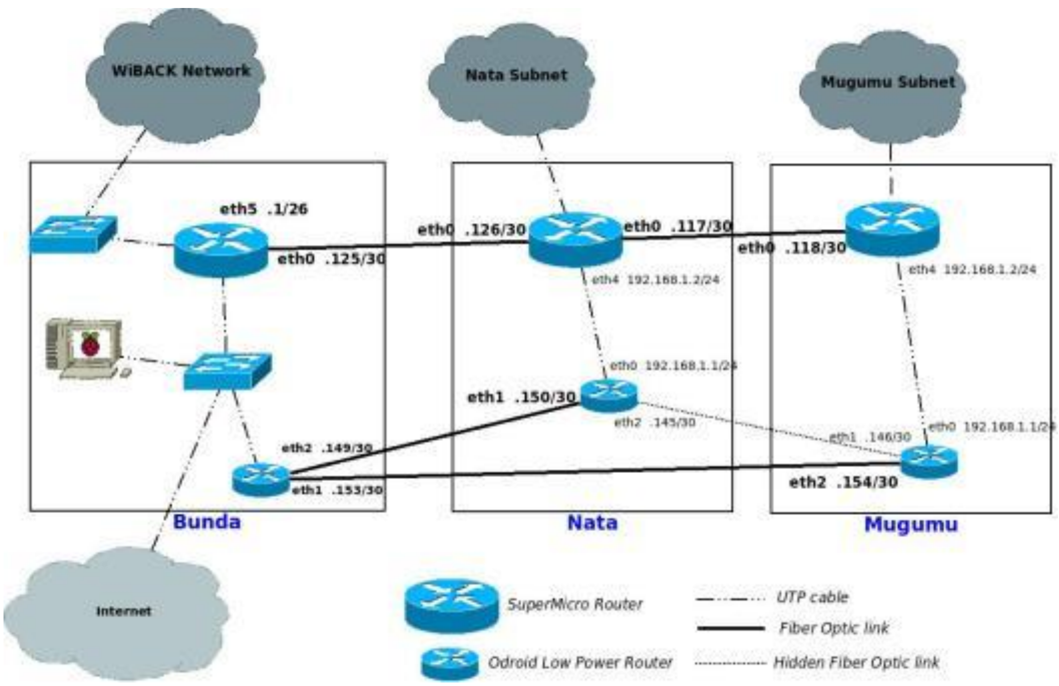

Figure 4: Network Topology of SBN, 2014-10-31

Three new 5 Watt routers have been deployed in SBN network along with the former ATOM-based routers, a network plan and topology can be found in Figure 4. Two routers at Nata and Mugumu are deployed to run over solar energy only, in order to experiment the power supply system with minimum load. Benefit from higher transmission power of duplex SPF, we bridges Bunda and Mugumu directly via a fiber trend bypassing Nata. By forming a star-like topology instead of line, we expect a more robust backbone network without single-point-of-failure.

5. Conclusions

By accepting a reduction of the bandwidth of the network interfaces from Gigabit Ethernet (1000Mbps) to Fast Ethernet (100Mbps) caused by the USB2 bus bottleneck, we have been able to reduce the average power consumption of our backbone router from the 20W range to below 5W. This reduction in the demand for power has made the power supply considerably more manageable and facilitates the use of alternative power storage technologies such as ultra-capacitors. The reduction in size, weight and cost of the routers is also essential in rural settings.

Due to the still moderate traffic volumes in the Serengeti Broadband Network, where the routers are deployed, the reduction in bandwidth has so far not had an impact on the quality of service experienced by the users. We regard it as a matter of time until the performance bottleneck caused by the USB2 bus can removed, which will most likely facilitate increase the energy/bit ratio an order of magnitude allowing 1Gbps links.

The authors are currently in the process of evaluating a successor to odroid U3, the XU3 (Odroid U3/XU3) which supports USB 3.0 with a maximum bandwidth of 5 Gbit/s, and is also looking into the design of low-power USB3-based 1Gbps optical interfaces.

References

Alix http://www.pcengines.ch/alix.htm

Battery management and Standards: https://energypedia.info/wiki/Battery_in_a_Photovoltaic_Power_Supply_System_Standards, http://batteryuniversity.com/learn/article/how_heat_and_harsh_loading_reduces_battery_life

BeagleBone Black http://beagleboard.org/black

http://usb.org

Interface masters Niagara 82048, http://www.interfacemasters.com

Iperf Performance Testing Tool https://fasterdata.es.net/performance-testing/network-troubleshooting-tools/iperf-and-iperf3/

Jager, Tjapko. (1982). *The Soils of the Serengeti Woodlands, Tanzania*. Wageningen: Centre for Agricultural Publishing and Documentation.

LMBench Tools for Performance Evaluation http://www.bitmover.com/lmbench/

Lyconsys Fibergecko100 http://www.lyconsys.com/download/datasheet_fibergecko100_eng.pdf

Nungu, A. (2011) *'Towards Sustainable Broadband Communication in Under-served Areas: A Case Study from Tanzania'* PhD Thesis. KTH http://people.kth.se/~amnungu/publications.html

Nungu, A., Olsson, R., Pehrson, Jiawei Kang, Daniel Kifetew, D., and Rustamov, A. (2012b) "Inclusive Ubiquitous Access - A Status Report", IEEE Africomm 2012, Yaounde, Cameroon, November 2012

Nungu, A., Olsson, R., Pehrson, B. (2012a) 'Design and Implementation of an Inclusive Ubiquitous Access', *Int. Journal of Wireless Personal Communications*, Springer, November

2012

Odroid U3/XU3 http://hardkernel.com/main/products/prdt_info.php

Odroid XU3 http://hardkernel.com/main/products/prdt_info.php?g_code=G140448267127

Raspberry Pi http://www.raspberrypi.org/

Supermicro X7SPA
http://www.supermicro.com/products/motherboard/ATOM/ICH9/X7SPA.cfm?typ=H&IPMI=Y

Biographies

Dr. Amos Nungu is a Lecturer, researcher, and consultant at the Dar Es Salaam Institute of Technology (DIT) in Dar es Salaam Tanzania. His research interest includes ICT for Development (ICT4D) in general, especially on establishment of sustainable broadband markets in under-served areas, looking at business models, services and technical solutions. Amos holds a PhD and MSc degrees in telecommunication systems from KTH Royal Institutes of Technology, Sweden. Dr. Nungu holds various managerial and professional responsibilities. Currently he is the Director for India-Tanzania Centre of Excellence in ICT. He is also the Executive Secretary for Tanzania Education and Research Networks (TERNET).

Robert Olsson, Senior Research Engineer, BSc from Uppsala University. Pioneer in high-speed open source routing with several contributions to Linux kernel networking code including the ipv4 routing algorithm. He is the founder of Bifrost/Linux distribution mainly intended for network infrastructure. Current work covers low-power networking, micro-controller software and hardware for renewable energy applications and sensor networks. Current employers KTH/NSLAB and Uppsala University. He also is the founder of Radio Sensors AB a R&D start-up in IoT, WSN and new energy.

Jiannan Guo is a master student of EIT ICT Master School in KTH from 2013 to 2014, majoring in Communication System Design. He was previously studying Internet Technology and Architecture at Technische Universität Berlin and Electronic Engineering at Fudan University of China. He participated in junior engineer training program of SPIDER and did internship in Tanzania in 2014.

Björn Pehrson is a professor emeritus in telecommunication systems affiliated with KTH since 1987. He was previously with Uppsala University and Swedish Institute of Computer Science. His research interests include how to establish broadband networks in under-served areas and Internet of Things, especially sensor networks for environment monitoring