Research Directions in Energy Efficient Networking

Raimo Kantola, Jörg Ott, Jukka Manner
{firstname.lastname}.tkk.fi
Department of Communications and Networking
Aalto University, Finland

Introduction

Energy consumption of the mobile and core packet networks is growing alarmingly if the current trend of growth of traffic volumes continues [1]. It is time to look for approaches of meeting the demand in traffic growth while optimizing for energy efficiency. In this paper, we make an attempt to classify the possible approaches, discuss the tradeoffs between and in the approaches, define research directions and point out some challenges and potential gains.

The premise of the classification is the result that the higher up in the protocol stack switching takes place the more energy is consumed [2]. We classify the approaches to (1) optical networking, (2) electronic networking. The latter we break down to (2.1) new architectures and (2.2) incremental improvements. Under new architectures we discuss energy efficiency in (a) addressing, (b) access, (c) routing and traffic engineering, (d) forwarding and (e) caching. Under incremental approaches we identify link level and nodal sleep modes and other node technologies. Incremental solutions are emerging on the markets in the form of Energy Efficient Ethernet and vendor specific solutions for their nodes.

Optical and Electronic

Optical transmission is energy efficient because the transmitted power is tightly focused and dispersion is minimal. Compare this to a base station with a spherical beam form. Optical circuit switching is also power efficient. However, circuit switching is unable to allocate network capacity efficiently to bursty traffic with short flows that are typical in the Internet. We can see this as a tradeoff between inefficient and slow allocation of network resources to traffic demand leading to wasting network capacity and power efficient transmission and switching on the optical layer. Other forms of optical switching are not mature yet. Any form of optical switching can trade off capacity allocation efficiency for the gain in power consumption. This trade-off can take different forms. An incremental solution is a hybrid or multimode switching node that can alternatively handle traffic in the optical or the electronic domain.

New architectures

Flat *addressing* in the sense that all devices have an address in the same address space is not energy efficient. The justification of long addresses is that many devices or device groups need different treatment in wide area routing. This implies large routing tables and a lot of power hungry fast memory. In [3, 5], we propose to use *globally unique names and locally significant identities and locally significant addresses* for global communication. Our proposal is one among the ID/address separation and edge based tunneling architectures. LISP falls into the same category. By removing host network ASes and host network address prefixes from the core routing tables, significantly simpler and power efficient core nodes become feasible. Addresses are one of the commons in networks. The longer addresses are carried in each packet the more address space is wasted by users. If routing locators are separated from host/service/user identities, core routing tables will need to have entries only for routing locator prefixes if the routing tables remain in the core at all. When allocated reasonably, rlocs can be aggregated and ISP network topologies can be hidden from each other leading to smaller FIBs. The size of the core FIB with entries for rlocs prefixes only is likely to be less than 10 000 entries for a non-default routing core for the present Internet. This assumes that core network operators prefer large nodes over small ones. This is a significant improvement from the present Internet core of around 300 000 routing entries. Considering core router FIB memory consumption, equally important than the Internet FIB is the use of memory for virtual networks.

Access can be interrupt driven or based on polling like in the present 3G and LTE architecture using UNSAF. Polling leads to battery exhaustion but also adds traffic into the access network. Introducing sleep mode into the access nodes becomes complicated because of polling: the sleep mode manager would need to know how to ignore the polling traffic in decision making. We have proposed a solution for an interrupt driven access architecture [3].

Recently, approaches that do not need *routing* tables in the core have been proposed using either bloom filters or link coloring [4]. These can be seen as forms of source routing. A shortest path routing algorithm uses a link state data base that is populated by a network discovery protocol. Creating a routing table for the next hop from the link state database reveals the whole path to the egress node. Therefore, source routing eliminates repetition of almost the same calculation in the core nodes saving energy while not complicating the edge too much. It is possible to pack source routes efficiently minimizing header overhead and thus saving network capacity and consequently energy. Source routing also leads to very simple core nodes: all the routing complexity is pushed to the network edge. One operator does not wish to reveal a source route showing its network topology to another operator. This limitation can be overcome.

On the edges, due to lower processor clock rates, the complex routing functionality can be executed with minimal energy consumption. The core nodes need to provide very high capacity for desirable network topologies to become feasible. These nodes will use very high clock rates. Therefore, simplifying the functionality of the core nodes saves energy. Simple here means that as few transistors are used for processing each packet as possible.

Source routing can be designed to support route pin-down. This facilitates easier control of the use of network resources and traffic engineering for minimizing power consumption. TE in the interest of maximum capacity leads to keeping all links evenly occupied. This renders link level sleep mode inefficient. Energy efficient routing does the opposite: it must concentrate traffic onto the minimal cut of links that allow meeting delay and capacity constraints. This enables the sleep mode on a maximum set of links and nodal interfaces connected to those links. Route pin-down helps to move traffic from lightly loaded links to the serving links that the TE system has defined. As a result, more links can go to sleep mode. This requires that the ingress node that populates the packets with the source route, besides the primary route, must store alternative routes as well.

Caching of popular content reduces core network traffic and cuts delays perceived by the user. This may give a significant power saving because power consumption of the core grows fast as a function of traffic volume. If a piece of content in a cache is never used after the storing, the network wastes memory and a minimal amount of energy. There is an optimum popularity rating of a content object such that all objects whose rating is higher should be cached in order to minimize power consumption. Costs of caching are compensated to the serving ISP by lower interconnection and transit charges towards tier 2 and tier 1 ISPs.

Finally, we come back to the premise: the lower in the stack switching takes place the less energy is consumed. Over the past 15 years the Internet architecture has deteriorated. More and more of the traffic is switched on application layer mainly because of the missing trust architecture. In the light of our premise, this is the worst architecture for power consumption in a packet network. The application level boxes are firewalls, session border controllers and application level gateways. Their purpose is network hiding or protecting user's and networks from harmful traffic. This speaks for a new network architecture that is able to provide network hiding on the Ethernet layer. This calls for topology driven allocation of MAC addresses and using Ethernet forwarding in the core. It is a popular claim that Ethernet does not scale to global networking. We have shown [6] that Ethernet scales to global packet transport. We believe that by adding a limited set of functions, we can use Ethernet forwarding/switching in Internet scale.

Incremental solutions

It is possible to save energy in data centers and network nodes for example by novel nodal power solutions. We are engaged in projects in this direction. This is important for being able to assess the value of any new architecture.

Research directions for EU-Japan collaboration

Enhancing Ethernet for Internet like networks, energy efficient routing and forwarding, locator/identifier split, trust architecture for the Internet, scaling network link speeds as a function of traffic.

Conclusions

We have tried to identify an approximate network structure that will give the minimum of energy consumption at a traffic that we expect in the next 10 years. Essentially, the conclusion is that switching, if not optical should be on Ethernet layer as often as possible and the highest capacity nodes should be simplest and use the minimum number of transistors/switching a data-unit. This is because the higher is the node capacity the higher clock speed for processing and power consumption grows relative to the square of the clock rate. The role of electronics switching is to adapt the network to bursty traffic consisting of short flows while power efficient optical circuit switching can be leveraged to providing VPNs. Additionally, low capacity VPNs must be supported on the Ethernet layer. Tunneling based edge can de-couple the technology choices of operators from corporations and core operators from access operators although some of the proposed techniques try to avoid this de-coupling. The de-coupling is a necessary precondition for the deployment of any large scale architectural changes into the network in the interest of power saving. Each operator should be free to do technology selection independently. While we need to do fundamental research on the new architectures, each operator can easily deploy incremental node level power saving solutions. The downside of both link and node level sleep modes is lower responsiveness to traffic changes. Gaining efficiency from link level sleep modes requires attention to power efficient routing and TE.

References

[1] S. Namiki, et.al, "Dynamic Optical Path Switching for Ultra-Low Energy Consumption and Its Enabling Device Technologies," saint, pp.393-396, 2008 International Symposium on Applications and the Internet, 2008

[2] G. Eilenberger, Energy efficient and scalable packet transport networks, Proceedings of the 2nd Japan-EU Symposium on the Future Internet, Tokyo, 13-14 October, 2009.

[3] R. Kantola, Implementing Trust to Trust with Customer Edge Switching, Perth, AMCA WS at AINA, 2010.

[4] V. Holopainen, R. Kantola, T. Taira, O-P. Lamminen: *Automatic Link Numbering and Source Routed Multicast*, in Proc. of AIMS, 2010.

[5] www.re2ee.org

[6] http://www.ict-etna.eu/