

Optimisation of Power Consumption in Wired Packet Networks

Erol Gelenbe¹ and Simone Silvestri²

¹ Dept. of Electrical & Electronic Engineering, Imperial College, UK
`e.gelenbe@imperial.ac.uk`

² Department of Computer Science, Sapienza University of Rome, Italy
`simone.silvestri@di.uniroma1.it`

Abstract. Over 500 million host computers, three billion PCs and mobile devices consume over a billion kilowatts of electricity. As part of this “system” computer networks consume an increasing amount of energy, and help reduce energy expenditure from other sources through E-Work, E-Commerce and E-Learning. Traditionally, network design seeks to minimise network cost and maximise quality of service (QoS). This paper examines some approaches for dynamically managing wired packet networks to minimise energy consumption while meeting users’ QoS needs, by automatically turning link drivers and/or routers on/off in response to changes in network load.

1 Introduction

Electrical energy is needed for Information and Communication Technologies (ICT) both to operate and cool the equipment. However ICT plays a complex role in energy consumption: as summarised by Francesco Serafini of HP: “ICT covers 2% of the global energy consumption. We can work to halve it to 1%, but the priority is to work to decrease the remaining 98%” (through greater use of ICT), via the “communicate more and travel less” paradigm. (cf. T. Cowen of BT): “if people rely heavily on videoconferencing, distance working, e-learning and e-commerce, there can be a sharp decrease in flights and journeys which will save money and carbon”. Thus computer-communication networks offer a way forward for reducing the consumption of energy and carbon emission by reducing land and air transport, but this potential reduction is partially offset by the power used by data centres and computer networks [1]. Network and web service providers have electrical annual costs in the billions of pounds. Even a fraction of energy savings in networks could lead to reduced financial costs and carbon savings. Therefore *future computer networks should include energy aware traffic management and routing techniques*, together with efficient hardware level energy management, to provide acceptable levels of QoS, at the lowest possible energy levels. The importance of *packet networks* on energy consumption increases with the “convergence to the Internet Protocol (IP)” paradigm whereby most modes of communication, including mobile telephony, are increasingly supported by underlying packet networks. Since IP networks rely on network nodes

and links, the electrical energy used for operating and cooling the equipment creates a crucial need for research on energy saving strategies for networks.

Our research starts from the premise that if one knows (a) the traffic being carried by a network, (b) the Quality of Service (QoS) requirements of the traffic flows, and (c) the capacity and energy requirements of the nodes and links that are available to store and forward this traffic, then (d) one could rationally select a set of nodes and inter-node links, and paths through these selected nodes from source to destination, so as to satisfy the required QoS at a minimum energy utilisation. Although much work has been done for power management in wireless networks, this subject has not yet received much attention for wired networks. In the wired case, the problem we describe can in principle be formulated as a standard network topology optimisation problem where the cost to be minimised is power consumption, and the constraint to be respected is QoS. However a *static optimisation* approach, assuming a given fixed set of traffic flows, is impractical due to dynamic changes in traffic levels. Approaches that can take advantage of short period changes in traffic levels, e.g., hourly or even shorter, would be more useful, and one can also consider variations in electric energy costs. More sophisticated approaches, which we will not study in this project) might consider the impact on carbon emissions. Moving data centres to remote locations where energy is cheap and outside temperatures are low (for cooling) has been considered, but this idea is impractical for networks where much of the infrastructure should be close to the end users.

Our work considers a *dynamic approach* to energy optimisation in packet networks, where link drivers and/or nodes are turned on or off, in response to traffic load in the network, with ensuring changes in the paths followed by the traffic so as to meet the QoS needs of the flows. Since it is technically easier to keep routers “on” constantly, and just turn the links’ drivers on/off according to need, and since drivers can use up to 40% or more of the energy consumed by a network node [2, 3], we will focus in particular on that approach.

Our main objective is therefore to develop a principled approach for the design of a software based “Energy Management System” (EMS) for wired packet networks that will make and carry out dynamic decisions to minimise network energy consumption while respecting QoS constraints. The EMS would run on top of the network (e.g., IP) layer, and *continuously* undertake the following steps:

- **a)** Observe ongoing traffic flows in the network, monitor the status of nodes and the network’s power consumption.
- **b)** Select the network configuration that offers an acceptable or better level of QoS to ongoing and predicted flows, with *lower energy consumption*.
- **c)** Manage and sequence dynamic changes in links and nodes, and reroute traffic, to achieve reduced power consumption at acceptable QoS levels.

This study does not address optimisation related to cooling systems for network routers because it would require knowledge of the physical spaces and external conditions related to buildings where routers are housed. However, when network

equipment consumes less power, there is an induced effect on heat dissipation power used for cooling.

1.1 Prior Work on Power in Wired Networks

Energy saving strategies for wired networks have not received much attention. Some preliminary works can be found in [2, 3, 4, 5, 6, 7]. Paper [4] characterizes the variation due to fluctuating electricity prices and argue that existing distributed systems should be able to exploit this variation for significant economic gains. In [5] some methods that allow for detection of inactivity periods in Ethernet networks are proposed. Such periods are exploited to obtain energy savings with little impact on loss or delay. Papers [2, 6] study the power profile of modern network devices and propose some method in order to reduce the energy consumption without affecting the performance. The authors of [3] considers jobs reallocation in order to reduce the number of active switches into the network. Finally, paper [7] surveys methods and technologies currently used for energy-efficient operation of computer hardware and network infrastructure in the context of cloud computing.

None of the previous works consider a general framework for energy optimisation in packet networks, where link drivers and/or nodes are turned on or off, in response to traffic load in the network, with ensuring changes in the paths followed by the traffic so as to meet the QoS needs of the flows. We introduced this problem in [8].

1.2 Prior Work on Power in Wireless Networks

Differently from the wired case, energy-saving routing protocols in wireless networks have been studied in detail. Energy saving techniques for wireless sensor networks [9], are important because of the limited power availability in networks which operate with batteries or renewable energy sources. “Topology control” (TC) algorithms [10, 11] for wireless networks have been proposed so as to dynamically modify the network graph to maintain or optimize desirable global properties, such as network capacity or user perceived QoS, while reducing energy consumption and wireless interference between nodes. These approaches dynamically adjust the transmission power of each node in order to save energy while guaranteeing connectivity. Another important criterion considered is to limit the ratio of the number of hops traversed by packets from the sources to the destinations, for a given power setting, to the number of hops traversed if all nodes were transmitting at maximum power, and some complex trade-offs occur as the nodes’ transmission power levels are varied. While the maximum power can yield the minimum number of hops, higher power levels will adversely affect collisions and interference, lower levels of transmission power potentially lengthen the paths traversed by packets, but reduce collisions and interference on each hop. As the hop count increases, the energy used per packet can also increase and adverse effects such as delay and loss can also increase.

In [12] topology control is studied with the addition of QoS constraints, and an integer linear programming problem is formulated so as to assign transmission radii to sensors (hence establishing the network topology) so as to minimise the overall global energy consumption while respecting the desired QoS constraints. Although it has similarities to the research we propose, the TC problem for wireless networks cannot be mapped directly into the problem we consider. In a wireless sensor network, a node may change the set of nodes with which it communicates in one hop by adjusting its transmission power, and hence range. This is not the case in wired networks, where a router is connected in one hop to routers which are turned on and connected to it via a link. In a wired node, power consumption may change as a function of the node's workload based on turning some or all the cores in the processors on an off, but this capability may be difficult to control explicitly, and in any case close to 60% of the router's energy consumption typically results from peripheral hardware. In TC for wireless nodes, one cannot switch off some wireless nodes because they may also act as sensors whose role is to gather information, in addition to forwarding packets.

Turning nodes off has been considered in [13] for an energy aware distributed data-centric routing protocol that takes into account the physical location of the sources of data; this is less relevant to the wired networks because wired packet networks may have source and destination nodes anywhere throughout the network. Finally, the broadcast nature of the wireless medium introduces both a degree of freedom (local multicast to neighbours) and additional constraints (interference and collisions) which do not exist in wired packet networks.

In [9] a middleware that manages the state (idle, active, sleeping) of sensor nodes so as to reduce the network's energy consumption, is proposed and two algorithms are considered, either turning off a node when the node is not involved in sending, forwarding or receiving data, or using information about the spatial "density of active nodes" so as to change the proportion of time spent in the active state (duty cycle) of nodes in proportion to the density of active nodes. In a sensor network the volume of traffic is generally relatively low and each node can alternate between different states (idle, active, sleep) quite frequently. This may not be possible in a wired scenario, where traffic volumes and speeds can be high, QoS constraints can be stringent, and large packet loss rates cannot be tolerated. Also, large wired routers can take a non negligible time to be turned on/off, especially when precautions must be taken to avoid packet losses in the queues. In [14] battery powered wireless ad-hoc networks are considered, and a routing technique which selects paths so as to both satisfy QoS constraints and minimises power consumption is suggested, so and some of these ideas may be used in our proposed research.

1.3 Power Minimisation in Clusters of Servers

Recent research has also considered energy savings in clusters of database and web servers, where one seeks to minimise power while guaranteeing acceptable levels of throughput and response time [15]. In such systems, energy consumption depends on CPU utilisation, but is also caused by components such as disks,

memory, network devices, etc., and an idle server may still use up to 60% of its peak power [15, 16]. Thus, in order to minimise energy power, it is necessary to turn off or down complete ICT servers as a function of load. In [16] policies based on economic criteria allocate resources within a large cluster of servers, focusing on energy, by managing the request dispatcher so as to concentrate the incoming load on a minimal set of active servers that fulfill the QoS agreement, while other servers are maintained in a low-power idle state. The related dynamic provisioning problem is studied in [17] for long-lived TPC connections, as in the case of instant messaging (Microsoft Messenger, Skype, etc.) or online gaming, using an approach that consists of a dynamic provisioning algorithm and a load dispatching policy. In [18], a dynamic provisioning technique is proposed for a platform that hosts several applications, using a queuing model for system analysis, and a provisioning technique. The analytical model yields the minimum number of servers required to respect the QoS agreement, and the provisioning technique uses a predictive strategy, based on long time scales, to determine future load. The prediction is then adjusted uses a reactive provisioning policy at small time scales, to address problems that may be caused by sudden changes in system load.

The main difference with the networked context is that the research on data centres addresses the question “What is the least number of processing nodes that are needed?”, while in the packet network context we first identify “*Which* nodes are needed within the given topology, and then decide how traffic should be re-routed through the active nodes. The networking context is also more complex to manage in the presence of ongoing traffic flows that will need to be re-routed so as to constantly satisfy the QoS constraints.

2 Technical Approach

As indicated earlier, this work focuses on dynamic energy optimization in the context of network routing, with monitoring of the current flows and prediction of future flows in the network, including source and destination pairs, volume of traffic in each flow, QoS constraints per flow, the path for each flow, and different “activity levels” ranging from “sleep” to “fully active” that each of the router nodes can have. Our work therefore addresses:

- (A) The formalisation of dynamic network management for energy optimization, and development of the optimisation algorithms, and
- (B) The design and implementation of the Energy Management System middleware (EMS) and its experimental evaluation in our network test-bed.

2.1 Dynamic Network Management for Energy Optimisation

For (A), we suggested two approaches: (i) the first assumes that the drivers of the links from a node can be turned on or off, resulting in proportional (possibly non-linear) changes in the energy consumption, (ii) the second, less realistic

approach with today’s technology, considers turning on and off certain nodes, with larger changes in energy consumption but at a slower rate than (i), say in the minutes. We think that (i) is more realistic because the node itself remains “awake” but its communication drivers are turned on/off.

Both approaches can be studied using a representation of the network N as a set of bi-directional links and nodes forming a graph. A link (i, j) $i, j \in N$, has a maximum traffic carrying capacity of $C(i, j)$ packets/seconds. At time t , link (i, j) can be in states: $k(i, j, t) \in \{0, 1, \dots, M\}$ where 0 is the “off” state, and M is the state in which the link operates at maximum capacity $C(i, j)$; in practice we can expect that there will be just two such states, 0 and M . In the case (i) a link’s traffic carrying capacity at time t is

$$K(i, j, t) = k(i, j, t)C(i, j)/M$$

For each $k(i, j, t)$ a link will have a power consumption (energy per unit time) of $P(k(i, j, t))$ while a node has a power consumption denoted by $P(i, t)$. Clearly for symmetric links we have $k(i, j, t) = k(j, i, t)$. In the case (ii) the *node* will have just two states, on or off, and we will be dealing at each instant of time with a sub-network $n(t) \subseteq N$, and the network $n(t)$ has an instantaneous power consumption of

$$P(n(t), t) = \sum_{i \in n(t)} P(i, t) + \sum_{i, j \in n(t)} P(k(i, j, t))$$

Let $k(t)$ be the matrix $[k(i, j, t)]_{n \times n}$. At time t , the network carries a set $F(t)$ of flows, each characterized by a source-destination pair, and each flow $f(t) \in F(t)$ has a traffic volume $V(f(t))$ and is of type $T(f(t))$ (e.g., UDP, TCP etc.), with *QoS constraints* $Q(f(t))$ which is typically a vector of numerical values (e.g., delay, loss, jitter) which must not be exceeded if the QoS is to be satisfied. The routing scheme at time t is $R(F(t), n(t), k(t), t)$, which assigns each flow to a sequence of nodes in $n(t)$ (for (i) $n(t) = N$) going from the flow’s source to its destination. Since flows interact with each other via the routing scheme, the *observed QoS* for flow $f(t)$ is $q(f(t), R(F(t), n(t), k(t), t))$, a vector denoting the *observed* or measured QoS for $f(t)$. If the network respects the QoS constraints at time t , then

$$q(f(t), R(F(t), n(t), k(t), t)) \leq Q(f(t)) \quad \forall f(t) \in F(t)$$

meaning that, say, the observed delay, loss and jitter for each flow are smaller than the required upper bound for delay, loss and jitter. The power optimisation problems can then be formulated as follows:

- **For approach (i):** Given a set of flows $F(t)$, find a link-level activity matrix $k(t)$ and an assignment of paths to flows $R(F(t), N, k(t), t)$, such that the QoS constraints are respected:

$$q(f(t), R(F(t), N, k(t), t)) \leq Q(f(t)) \quad \forall f(t) \in F(t)$$

and power consumption $P(N, t)$ is minimised.

- **For approach (ii):** Given a set of flows $F(t)$, find a sub-network $n(t) \subseteq N$ and an assignment of flows to paths in $n(t)$, $R(F(t), n(t), k(t), t)$, such that the QoS constraints are respected:

$$q(f(t), R(F(t), n(t), k(t), t)) \leq Q(f(t)) \quad \forall f(t) \in F(t)$$

and power consumption $P(n(t), t)$ is minimised.

To describe the optimisation problems, we can construct a “network of queues” model [19], where in (i), the link “service rates” are proportional to $k(i, j, t)$, so that algorithms can be designed to choose the matrix $k(t)$ as a function of the flows’ characteristics and QoS constraints so as to minimise power consumption. A differentiable cost function will be developed to seek local minima with gradient-type techniques to seek local minima, with the intention of deriving expect fast (polynomial time in the size of the network) optimisation algorithms using our experience with gradient descent learning in neural networks [20]. For (ii), to avoid the combinatorial explosion in seeking the best choice among all subsets of active nodes, we will examine greedy heuristics to identify the links and nodes that carry the least traffic, or the links and nodes carrying the least traffic among those with the highest power consumption, for possibly being turned off, and examine incremental policies for turning links or node off or on. This raises a challenging “real-time” aspect further discussed below.

As we change the network topology from network $n_1(t)$ to another one $n_2(t')$, $t' > t$, these networks will necessarily have common nodes (to avoid losing traffic), and hence traffic re-routing must accompany changes in network topology. Thus re-routing must be carried out in stages, so that once the future sub-network $n_2(t')$ is selected then:

- First certain flows should be diverted to nodes which are common to the current network and the future network, to avoid traffic loss during re-routing,
- Then certain links or nodes are turned off, while others may be turned on, so that the new network is established.
- Finally traffic has to be re-routed to take its new form.

These steps must be taken in a manner that the QoS for each flow is respected and the overall integrated power consumption over all the steps is lower. In networks operated by Internet service providers, excess capacity is usually maintained so as to protect the users’ traffic in the case of node and link failures of nodes and this feature actually simplifies the practical implementation of our proposal. We will therefore examine how this feature can be used to advantage in both (i) and (ii), as traffic is re-routed so as to turn off certain links or nodes to reduce power consumption.

2.2 Design of the EMS

The EMS, the software system which will implement the optimisation algorithms, should use a proactive approach to monitor the network and explore

control decisions. It will collect information about system state, then compute a course of action, and decide on the sequence of changes to be made. The EMS will gather information about: 1. *Active traffic flows*: sources and destinations, information about paths, traffic rates, experienced QoS (delay, jitter, packet loss), 2. *Power consumption*: instantaneous power consumption of each link driver and router, 3. *Router traffic*: identity (flow) and amount of traffic passing through each router.

The EMS uses the optimisation algorithms that we develop to evaluate whether a change in the network can provide significant reduction in power consumption with acceptable QoS, and evaluate whether the sequence of actions to reroute traffic and turn off and on certain (i) links (drivers) and/or (ii) routers can be carried out without major QoS degradation. Then it acts upon the routing tables to redirect traffic, then turn some drivers and nodes off, and it might turn others on, then again re-redirect traffic, while continuing to monitor the QoS and power consumption.

The role of the EMS bears some similarity with previous work on the Cognitive Packet Network (CPN) routing protocol [21] that monitors the QoS state of flows, and recommends the choice of paths to optimise users' QoS.

CPN is implemented in the Imperial College test-bed described in <http://san.ee.ic.ac.uk>, using “smart packets” (SP) which travel through the network from node to node to collect QoS information at nodes and links. The resulting information is then returned to the sources, using “acknowledgement packets” (ACKs) which differ from conventional ACKs in networks. In CPN, the traffic payload is carried by “dumb packets” which use source routing. In CPN decisions are distributed, and sources select the paths that they use. In EMS on the other hand, we will begin with a *centralised* approach, even though a distributed approach may be taken at a second stage. We will also consider how the EMS may combine “admission control” for new connections to power optimisation and traffic routing. However, once the EMS determines the links or nodes to be turned off, and chooses the new paths for the flows that are affected by this change, it can inform the source nodes of the flows and provide them directly with the new source routes before it actually makes the changes at the link or node level. Like CPN, EMS may also use SPs to (I) collect QoS information about nodes and paths, and (II) collect data about ongoing energy consumption at nodes. The latter may be done with sensors placed at the power plugs of drivers or nodes or from their power supplies. Using the same SP to collect both QoS data and power data would also provide better timing accuracy about the network state which is of interest and includes both power and QoS. Once the best options have been chosen and the sequence of decisions are selected, the EMS can also distribute the routing decisions throughout the network using a mechanism similar to MPLS, or use source routing as indicated earlier.

3 Some Preliminary Experiments

In order to illustrate the feasibility of our research, we ran some simple experiments on the test-bed of Figure 1 representing the SwitchLAN network topology

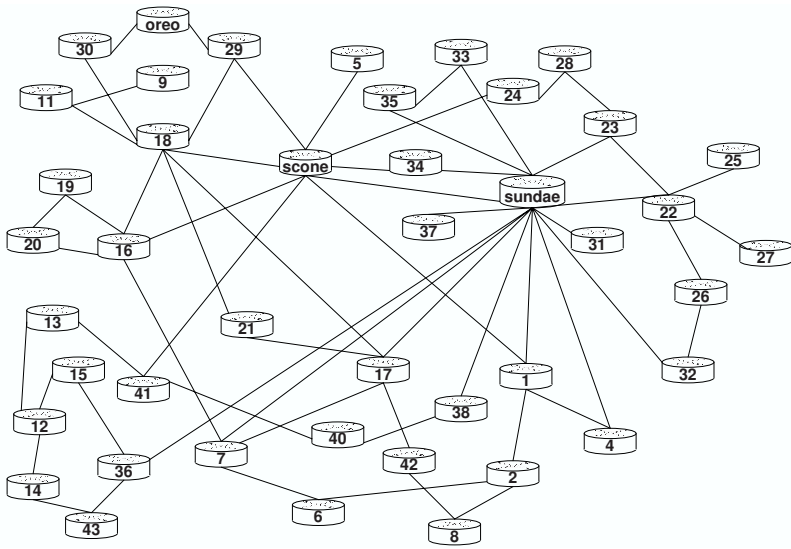


Fig. 1. Topology of the test-bed in use

[22]. The network is composed by 44 routers running the CPN adaptive routing protocol [21]. We emulated the effect of disabling the drivers some of the nodes' links by blocking their links.

The first experiment aims at studying QoS metrics of interest, namely delay, packet loss and jitter, while turning some routers off. We activated three 1Mb/s flows from node 8 to 1, 13 to 40, 30 to 21. Every 100 seconds we deactivated one router by disabling its links. In particular we deactivated nodes 28, 2, 34, 41, 16, 18, 26, emulating a 10-20% reduction in power consumption. Figures 2, 3 and 4 show the measured delay, packet loss and jitter, respectively.

These results highlight that not all topology changes affect the perceived QoS. In particular, changes made at instants 100s, 300s, 500s and 700s do not have a particular influence on the delay and packet loss, since the paths used by the routing protocol do not contain any of these nodes. By contrast, changes made at instants 200s, 400s, and 600s result in an increase of the measured delay and in momentary peaks of packet loss. Delay increases because these nodes lie on the shortest path among sources and destinations, thus longer paths have to be used at each change. The reason of the peaks in packet loss is instead twofold. On the one hand, some packets are still stored in the router queues at the moment in which links are disabled. On the other hand, the CPN protocol needs to detect the change and eventually discover alternative paths to the destination. Figure 4 shows that jitter remains unchanged in this experiment.

From these results it is possible to deduce that our research can effectively reduce the energy consumption, while guaranteeing an acceptable level of QoS, by turning on and off routers according to the network load.




Fig. 2. Delay - first experiment

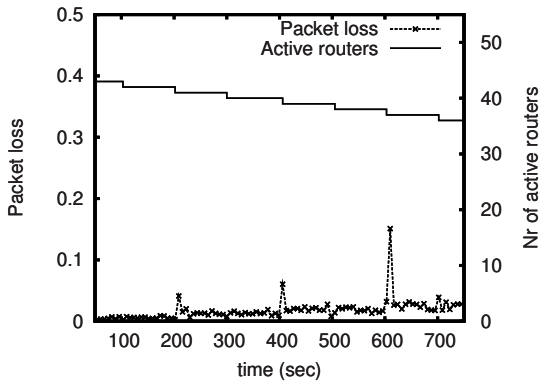


Fig. 3. Packet loss - first experiment

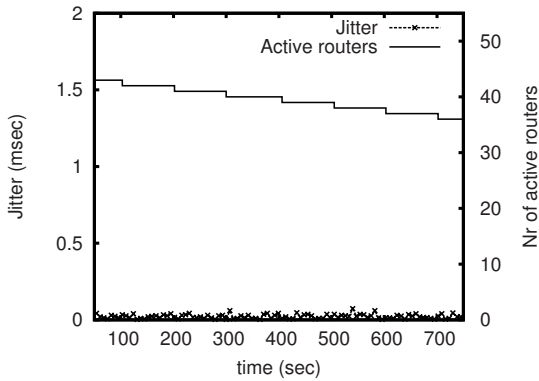


Fig. 4. Jitter - first experiment

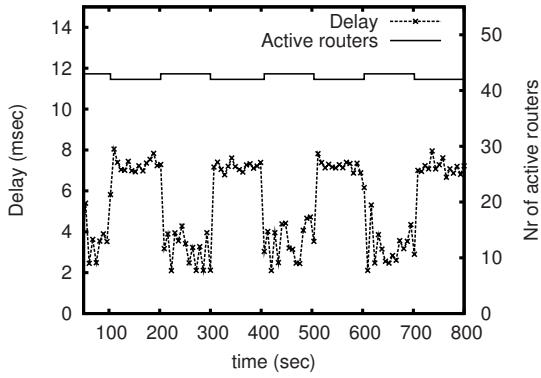


Fig. 5. Delay - second experiment




Fig. 6. Packet loss - second experiment

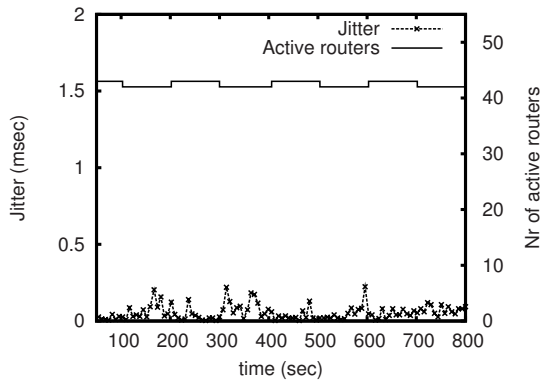


Fig. 7. Jitter - second experiment

The second experiment aims at showing the CPN protocol capabilities of reacting in consequence of a change in the network. We activate three 1Mb/seconds flows from node 23 to 38, 17 to 22, 32 to 33, and we turn *sundae* off (by disabling all of its links) and then on, every 100 seconds. Each flow has a shortest path through *sundae*, thus every time it is turned off the CPN protocol has to adapt the routing scheme to the new topology by using longer paths. Figures 5 and 6 show the measured delay and packet loss, respectively. The network runs CPN which adapts the paths so as to avoid *sundae* when it is off. Although delay and loss increase, they recover rapidly after each change, while jitter appears to remain stable.

4 Conclusions

This paper describes some preliminary results on research that can contribute to discovering principled techniques that can reduce energy consumption in packet networks. It is therefore hoped that this work can influence the ICT research community and industry to direct more attention to these issues.

As a consequence, energy optimisation can become a new and important area of network optimisation and management that directs researchers' attention to energy savings in computer networks.

References

1. Fan, X., Weber, W.-D., Barroso, L.A.: Power provisioning for a warehouse-sized computer. In: ISCA 2007: Proceedings of the 34th annual international symposium on Computer architecture, pp. 13–23. ACM, New York (2007)
2. Mahadevan, P., Sharma, P., Banerjee, S., Ranganathan, P.: Power awareness in network design and routing. In: Proceedings of the IEEE International Conference on Computer and Communication, Infocom (2008)
3. Mahadevan, P., Sharma, P., Banerjee, S., Ranganathan, P.: Energy aware network operations. In: Proceedings of the IEEE Global Internet Symposium (2009)
4. Qureshi, A., Weber, R., Balakrishnan, H., Gutttag, J., Maggs, B.: Cutting the electric bill for internet-scale systems. In: Proceedings of ACM Sigcomm (2009)
5. Gupta, M., Singh, S.: Using low-power modes for energy conservation in ethernet lans. In: Proceedings of the IEEE Infocom Minisymposium (2007)
6. Chabarek, J., Sommers, J., Barford, P., Estan, C., Tsiang, D., Wright, S.: Power awareness in network design and routing. In: Proceedings of the IEEE International Conference on Computer and Communication, Infocom (2008)
7. Berl, A., Gelenbe, E., di Girolamo, M., Giuliani, G., de Meer, H., Dang, M.Q., Pentikousis, K.: Energy-efficient cloud computing. *The Computer Journal* (August 2009)
8. Gelenbe, E., Silvestri, S.: Reducing power consumption in wired networks. In: Proc. 24th Annual Int. Symp. Computer and Information Sciences, ISCIS 2009 (2009)
9. Xu, Y., Heidemann, J., Estrin, D.: Adaptive energy-conserving routing for multi-hop ad hoc networks. USC/Information Sciences Institute, Research Report 527 (October 2000), <http://www.isi.edu/~johnh/PAPERS/Xu00a.html>

10. Rajaraman, R.: Topology control and routing in ad hoc networks: a survey. *SIGACT News* 33(2), 60–73 (2002)
11. Santi, P.: Topology control in wireless ad hoc and sensor networks. *ACM Comput. Surv.* 37(2), 164–194 (2005)
12. Jia, X., Li, D., Du, D.: Qos topology control in ad hoc wireless networks. In: *Proc. of IEEE INFOCOM 2004* (2004)
13. Boukerche, A., Cheng, X., Linus, J.: A performance evaluation of a novel energy-aware data-centric routing algorithm in wireless sensor networks, vol. 11(5), pp. 619–635. Kluwer Academic Publishers, Hingham (2005)
14. Gelenbe, E., Lent, R.: Power-aware ad hoc cognitive packet networks. *Ad Hoc Networks* 2(3) (2004)
15. Economou, D., Rivoire, S., Kozyrakis, C., Ranganathan, P.: Hardware-agnostic full-system power modeling. In: *MOBS* (2006)
16. Chase, J.S., Anderson, D.C., Thakar, P.N., Vahdat, A.M., Doyle, R.P.: Managing energy and server resources in hosting centers. In: *SOSP 2001: Proceedings of the eighteenth ACM symposium on Operating systems principles*, pp. 103–116. ACM, New York (2001)
17. Chen, G., He, W., Liu, J., Nath, S., Rigas, L., Xiao, L., Zhao, F.: Energy-aware server provisioning and load dispatching for connection-intensive internet services. In: *NSDI 2008: Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*, pp. 337–350. USENIX Association, Berkeley (2008)
18. Urgaonkar, B., Shenoy, P., Chandra, A., Goyal, P., Wood, T.: Agile dynamic provisioning of multi-tier internet applications. *ACM Transactions on Autonomous and Adaptive Systems (TAAS)* 3(1) (March 2008)
19. Gelenbe, E., Pujolle, G., Nelson, J.C.C.: *Introduction to queueing networks*. John Wiley Ltd., New York (2000)
20. Gelenbe, E., Hussain, K.: Learning in the multiple class random neural network. *IEEE Trans. on Neural Networks* 13(6), 1257–1267 (2002)
21. Gelenbe, E., Lent, R., Nunez, A.: Self-aware networks and QoS. *Proceedings of the IEEE* 92(9), 1478–1489 (2004)
22. The Swiss Education & Research Network, <http://www.switch.ch/network/>