Deriving a Dynamic Power Management Model for Next Generation Network Core Routers

H. Abaci, G. Parr, S. McClean, A. Moore, L. Krug (BT Adastral Park)
India-UK Advanced Technology Centre of Excellence in Next Generation Networks, Systems and Services
School of Computing & Engineering, University of Ulster, Coleraine, Northern Ireland, UK, BT52 1SA
(Abaci-H@email.ulster.ac.uk, gp.parr@ulster.ac.uk, si.mcclean@ulster.ac.uk, aa.moore@ulster.ac.uk, louise.burness@bt.com)

Keywords: **Power Model; Power Saving; NGN; Core: Metro.**

1. INTRODUCTION

Experts estimate that IP traffic will continue to grow at 43% per annum, doubling every 1.4 years [3]. Previous estimates show that ICT already consumes 2% - 3% of world energy consumption. Increasing network activity will lead to an increase in the energy required to move traffic around the network. Energy consumption, with related costs and CO2 emissions is a major concern, but there are also environmental effects associated with new equipment such as the extraction of materials (phosphorus, mercury) and carbon dioxide (CO_2) involved in manufacturing.

The GeSI study posits that around 460 Mt of CO2e emissions could be saved by 2020, by making use of telecommuting, videoconferencing, e-paper, e-commerce and online media [2]. Whilst these savings are greater than the additional costs on the network, there is still a need for the network to respond to the regulatory pressures (The Climate Change Act [1]), to reduce, rather than grow its energy use. Urgent optimisation of ICT infrastructure is required.

In this poster we analyse power consumption within the fabric of a Next Generation Network (Slide 2). This is a particularly difficult area to optimise as the performance – QoS and availability - of the metro-core networks is of paramount importance. However this does not imply that energy-saving techniques cannot be applied to such routers. To that end, we explore a typical Cisco CRS-1 core router in the field at the subcomponent level to highlight each state and utilization level and the workload impact on each component - namely, O/E (optical to electronic) conversion, buffer, forwarding engine, routing engine, switch fabric, power supply conversion and cooling. Our goal is to dynamically manage power within the metro-core routers of a network by determining energy efficient paths to send data taking into account performance and availability. This will require two significant developments: firstly (investigated here), we need a way to adapt the power of components to conditions. This is useful in its own right, as many components often run at very low levels of utilisation. Secondly, an enhancement to the routing system could select paths that would enable the energy of the entire network to be optimised.

2. SAVING TECHNIQUES

According to the literature, there are several ways to reduce energy consumption within an ICT infrastructure. However, we will investigate the simplest saving methods, namely sleeping and slowing for network communication. These have to be examined in the context of maintaining SLAs.

Sleeping is applicable to a component or subcomponent of a networked device while idle. These idle times can be between data packets (interpackets) and/or data communication (inter flow). However, this requires some modification to current protocol specifications.

Slowing saves energy where sleeping is not an option [4] by sufficiently slowing down the transmission rate or CPU clock to meet requirements. This can also be beneficial in environments with very high availability requirements.

3. POWER MODEL

In this section, we introduce a new model (Slide 6) to calculate network device power consumption according to configuration and real-time utilization. Configuration includes the number of “on” slots in the device and the number of “on” ports in the slot. In this case, the overall device utilization depends on the level of port utilization.

We allow for components to operate in different states (e.g. standby or half rate). Any model should to be scalable to incorporate more states in the future, so we include the flexibility to add and remove components that may vary from vendor to vendor. Also, the model must include details at the port level to investigate techniques that establish savings at the link level rather than for the whole device.

The Power model and its notations are shown in the poster (Slide 6). The first and second term of the formula $(\beta_a^{col}(1-\alpha_{slp}^{dev})+(\alpha_{slp}^{dev}*\beta_{slp}^{col}))$ calculate how much power is consumed by the device for cooling.

The third term of the formula (three summations) determines the number of ‘on slots’ and ‘on ports’. Also α values determine how long a particular port stays in the each state (utilization), and their workload at the each component. Each component’s power consumption is determined by a job’s complexity and CPU intensity. Each state has a specific set of jobs to perform within sub-components. A fully utilized CRS-1 core router (all slots and all ports connected and 100% utilized), consumes [9] 10KW power of which 7% relates to o/e conversion, 5% buffer, 32% forwarding engine, 10% switch fabric, 11% routing engine and 35% power conversion and cooling [5]. These percentages are simply represented by β values in the formula, and each state has its own impact on these subcomponents.

3.1. Power Calculation of the Core Router on the Field

Here, we present the network diagram shown in the poster (Slide 5). Based on this scenario, we use a power formula to find the total energy-savings achieved by using sleeping and slowing techniques. Results are generated from a MATLAB model (Slide 8).

3.1.1. Configuration

We assume that 7 fully meshed off-the-self core routers (Cisco CRS-1) are connected to each other by 10 Gbps DWDM fibre links (Slide 5). These core routers are high capability large routers that provide IP routing at about 0.5 Tb/s switching capacity and are typically located in highly populated cities. We also assume that 5 metro routers are connected to each core router. For the sake of availability and resilience we connect each metro router to two different core routers via 1Gbps fibre links. For the metro-core connection we employed a 3x4 port 1Gbps Line Card and one of them has 2 empty ports. For simplicity, we assume that the configuration of all core routers is identical.

3.1.2. Comparison of the Saving Techniques

We assume that each router has a set of information tables maintained within each router. These are named the Router Configuration Table and Workload Table (Slide 5). This information is gathered by management plane traffic using a set of calculations performed by the system or the router itself. The Router Configuration table shows the number of available slots (non empty slots) and which ports are connected by a physical link on that slot. The Router Configuration Table also contains operating bitrates, average utilization, predicted utilization at the port, and to where the port is connected. The Workload Table keeps each router's sub-components and their associated Workload Parameters (WP) in different states. Each port can be in a different state. This information helps to derive energy utilization metrics and subsequently establish dynamic power management using saving techniques for any autonomous network. Any such calculations have to be made in the context of constraints that may be imposed for access and availability.

Non Power Aware, Router: For simplicity we assume that each port has 3% utilization in the Backup (idle) state, and cooling WP is 0.35 (35%) and O/E conversion is 0.07 (7%), because utilization has little effect on the cooling and power conversion and O/E conversion [5][6].

Power Aware – Sleeping Router: We assume that appropriate hardware support in the operating system calculates the best sleeping time interval between neighbouring routers. The routers switching on/off time can be set to 1ms, with a dummy packet prior to data transmission to ensure no data loss during wake up. The routers can use buffer and burst strategies to create more sleeping time, and this sleeping time pattern is similar to WoA (pareto - Bursty) [7] (see slide 7). In addition, this sleeping time can be increased by link route aggregation. Listening and buffering at the port causes some power consumption at the O/E conversion. We assume that this WP is 0.03 and little consumption happens at the power conversion, which has a WP is 0.02. No power draws from the rest of the sub-components.

Power Aware – Sleep & Slow Router: Slowing saving is applicable to a non-sleeping port and device. It has been suggested that it is better to apply sleeping when the port has less than 30% utilization. When it is more than 30% then rate slowing is appropriate. We assume that the system determines the future (required) bit rate and adjusts the link rate according to this future prediction. There are 10 uniform bitrates between 1Gbps to 10Gbps and 1 between 1Gbps to 100Mbps to switch. The switching time is less than 1ms, so there is no significant delay and packet loss during switching [7]. To avoid rate oscillations the operational bitrate is set higher than the predicted rate. For example if predicted bitrates is 8Gbps, the router should operate at 10Gbps.

4. DYNAMIC POWER MANAGEMENT

This research aims to save power by transmitting data via paths of lesser power consumption within multipath environments (Slide 3), and (if feasible) – to create more opportunities for sleep saving by switching off unused links, line cards or sleeping an entire node. However, aggregating traffic into a less utilized link will increase buffer size, packet drop, latency, hardware usage (forwarding engine, route engine, buffering etc.) and retransmission [8]. Moreover we should be aware that the overhead on an alternative path should not cause packet drop, delay or retransmission, which will comprise the QoE.

5. CONCLUSION

Idle metro-core routers consume significant power if no intelligent and context-aware power saving technique is enabled. Such a situation comes about because of highly variable traffic loads, and the need for resilience and redundancy (over provisioning). Whilst certain ports may be busy it is always possible that other ports may be idle for long period of time. Based on our MATLAB calculations, idle router consumption is 3699 W and fully utilized router is 5125W by given configuration above (Slide 8). Results shows that sleeping is the most effective saving for a less utilized device up to 30% and slowing is beneficial when a device is utilized more that 30 % and Sleep & Slow saving beneficial under 20 % utilization. Up to 48 % saving is achievable for a modestly configured core router (about 70 Gbps in operation out of a maximum of 620 Gbps switching capacity) in idle state, 39% saving for 5% utilized, 26% saving for 10% and 14% saving for 15%. Future work will involve further effort on the derivation of energy models in the presence of constraints and will also explore opportunities for use of optimisation techniques.

6. REFERENCES

[1] Green ICT , Briefing Paper, January 2009, http://www.jisc.ac.uk/media/documents/publications/bpgreenictv1.pdf .

[2] GeSI, "ICT and Climate Change: The need for a SMART Alliance", http://www.gesi.org/LinkClick.aspx?fileticket=dTfIRJ%2faW3g%3d&tabid=37.

[3] Cisco Visual Networking Index – Forecast and Methodology, 2007-2012," 16 June 2008.

[4] C. Gunaratne, K. Christensen, B. Nordman, and S. Suen, "Reducing the Energy Consumption of Ethernet with Adaptive Link Rate (ALR)", IEEE Tran. on Computers, vol. 57, no. 4, pp. 448–461, (Apr 2008).

[5] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, R. S. Tucker, "Energy Consumption in Optical IP Networks", IEEE, OSA, (2009).

[6] J. Chabarek, J. Sommers, et al. "Power awareness in network design and routing", In INFOCOM (2008).

[7] S. Nedevschi, L. Popa, G. Iannaccone, S. Ratnasamy, and D. Wetherall, "Reducing Network Energy Consumption via Sleeping and Rate-Adaptation", USENIX Association Berkeley, CA, USA, (2008).

[8] L. Chiaraviglio, M. Mellia, F. Neri, "Reducing power consumption in backbone networks", ICC Dresden, Germany, (2009).

[9] Cisco CRS-1 Datasheet, http://www.cisco.com/en/US/prod/collateral/routers/ps5763/prod_brochure0900aecd800f8118.pdf