Energy-Efficient Master-Slave Edge-Router Upgrade Paths in Active Remote Nodes of Next-Generation Optical Access

Michael C. Parker[1], Richard Martin[1], Stuart D. Walker[1], Ward van Heddeghem[2], Bart Lannoo[2]
1. School of Computer Science & Electronic Engineering, University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, UK
{mcpark, rmarti, stuwal}@essex.ac.uk
2. Dept. of Information Technology (INTEC) of Ghent University – IBBT, Gaston Crommenlaan 8, 9050 Gent, Belgium
{ward.vanheddeghem, bart.lannoo}@intec.ugent.be

Abstract: Our design rules offers maximally energy-efficient Gb/s→Tb/s edge-router upgrade paths. One path assumes 10% average traffic intensity with 68% energy-efficiency gains over 5 upgrades, while 30% traffic load enables 45% energy-efficiency gains over 9 generations.
OCIS codes: (060.0060) Fiber optics and optical communications; (060.4256) Network optimization

1. Introduction

Optimisation of router power consumption is an important issue [1] as offered bandwidth continues to increase exponentially (Nielsen's Law [2]), despite empirical estimates [3] that router power consumption P tends to increase sub-linearly with capacity C, i.e. $P \propto C^{2/3}$. This is especially true for edge-routers located at active remote nodes (ARNs), e.g. for Active Optical Network (AON) architectures in next-generation optical access networks, which may be locally powered (or assisted) by small-medium scale renewable (solar, wind) energy sources. One issue with respect to energy efficiency is that the power consumption of routers does not generally scale well with the traffic load, or even the capacity gradually installed over a router's lifetime [8]. Hybrid equipment approaches [4] and master-slave configurations of paired routers [5] offer important energy efficiency savings [6]. Here, the master equipment (with a high power requirement P_{master}) deals with those (relatively few) periods of time with high traffic demand, whilst the slave equipment is optimally designed to cope with smaller traffic loads, up to a threshold level L_T, and to consume a significantly lower power P_{slave}. A master-slave configuration is most appropriate for telecoms equipment exhibiting a relatively inelastic energy consumption profile with respect to traffic demand, e.g. routers are known toconsume up to 90% of maximum power even in their idle states [7]. Fig.1 shows an edge-router in a master-slave configuration. Additional advantages of the master-slave configuration include the fact that switching-on and powering-down times can be relaxed, since only the (simple) switch at the input requires a very fast switching time. In addition, redundancy in the master-slave configuration offers potentially greater resilience (and reliability) during operation, assisted by relaxed router equipment technical specifications, i.e. slower switching-on and -off times (with overlapping operating times), so that the master-slave devices can be operated less "aggressively". For lower traffic intensities, e.g. a normalised traffic load of α=10%, substantial energy-efficiency savings of η>68% are theoretically possible [5]. However, for future system upgrades, the migration trajectory offers possibilities for energy-efficiency (OpEx) as well as cost (CapEx) optimization. E.g. when upgrading ARN throughput capacity, the larger master edge-router of a first generation (I) configuration can be re-used to become the slave router for a master-slave configuration in the second generation (II) architecture, as indicated in Fig.1(b).

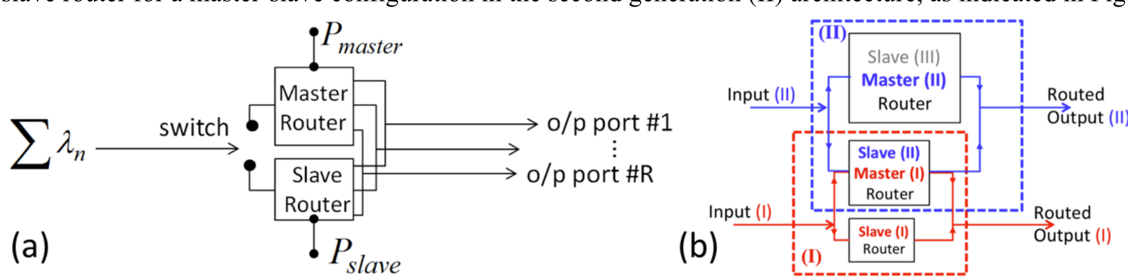

Fig.1. (a) Master-slave configuration for edge-routers; (b) Schematic of two successive generations (I) & (II) of master-slave configurations.

There is a trade-off between anticipated energy-efficiency gains as a function of expected average traffic intensity α, versus the CapEx costs of either adopting relatively few system upgrades by building-in generous future-proofing and hence greater over-provisioning; as opposed to assuming a higher frequency of system upgrades over time, but operating each system at closer to its maximum capacity (however with reduced opportunity for energy-efficiency savings). In the next section, we evaluate potential energy savings as a function of the traffic intensity.

2. Poisson Traffic Intensity Model

We assume that independent, identically Poisson-distributed packets arrive at the ARN via a Poisson process, such that for a maximum capacity N packets per second (i.e. corresponding to the master edge-router capacity, $C_{master} = \bar{b}N$, where $\bar{b}$ is the average packet length in bits), the instantaneous traffic load is n, with an average (modal) traffic intensity $L = \alpha N$, where α is the normalized average traffic intensity. The resulting probability that at any instant n packets per second are incoming is $p(n, L) = L^n e^{-L}/n!$, which is employed to be symmetrical about α=0.5 so that for a full capacity (α=1), as well as the zero load case (α=0), the distribution becomes a delta function. The total power dissipated by the master-slave configuration is given by the probabilistically-weighted sum of individual power consumptions of the two devices:

$$P_{total} = P_{slave} \int_0^{L_T} p(n)dn + P_{master} \int_{L_T}^{N} p(n)dn \quad (1)$$

Eqn.(1) indicates that once the traffic intensity reaches a threshold value L_T, then traffic is switched between the routers; the required value of L_T (which corresponds to the maximum capacity of the slave router $C_{slave} = \bar{b}L_T$) is an optimization problem, depending on the average (expected) traffic intensity α, and the relative values of P_{master} and P_{slave} (proportional to $C_{master}^{2/3}$ and $C_{slave}^{2/3}$, respectively.) The overall energy saving possible is given by the weighted summation (integration) of the operating powers of the two routers divided by the power of the master device:

$$\eta = 1 - \frac{P_{slave} \int_0^{L_T} p(n)dn + P_{master} \int_{L_T}^{N} p(n)dn}{P_{master} \int_0^{N} p(n)dn} \quad (2)$$

Fig.2 plots Eqn.(2) as the normalised threshold value $l_T = L_T/N$ is varied. As might be expected, for the two extreme cases of $l_T = 0$ and $l_T = 1$ when the master router is used for 100% of the time, then the marginal improvement (gain) in energy efficiency reduces to zero. Between these two extreme cases, however, the energy efficiency shows a maximum, particularly for the cases of low traffic loads, $\alpha \to 0$. Inset in Fig.2 is a plot of the locus of the energy-efficiency maxima for varying α, its convexity arising from the sub-linear $P \propto C^{2/3}$ relationship. Table 1 tabulates the values of the optimised maximum energy-efficiency gains η with their optimum normalised threshold levels l_T for different average loads α. As the average traffic intensity α increases towards maximum capacity, the possible energy-efficiency savings from the master-slave approach reduce greatly. Intuitively this can be understood by noting that the switch-over threshold l_T increases with α, meaning that the slave router capacity (and therefore its power consumption) approaches that of the master device. The normalised threshold load l_T can now also be understood to be the ratio of the master and slave router capacities: $l_T = C_{slave}/C_{master}$.

Table 1: Maximum energy-efficiency gains (h_{max}) at optimised normalized switch-over threshold (l_T=L_T/N) for different average traffic intensities α.

α	η_{max} (%)	Normalised Threshold Load (l_T) for Switching
0.1	68.1	0.18
0.2	54.1	0.30
0.3	45.7	0.40
0.4	34.3	0.51
0.5	26.3	0.61
0.65	16.9	0.74
0.8	8.3	0.86
0.9	2.8	0.94

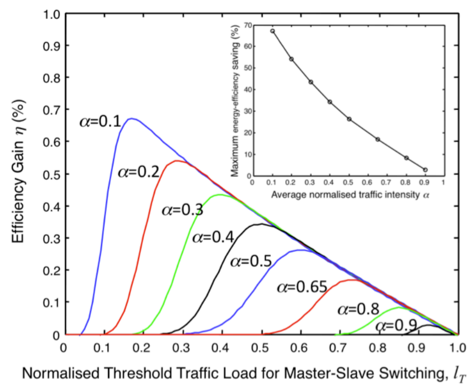

Fig. 2. Efficiency gain curves for different *a* values as a function of varying normalised threshold traffic load l_T

3. Future-Proofed Upgrade Trajectories

The optimized normalized threshold load l_T (which depends on the value assumed for α) can therefore be used as a design tool to calculate optimum values for master edge-routers in a future capacity upgrade scenario. Re-employing a master router as the new slave in a next-generation system means we adopt the following design rule:

$$C_{master}^{(II)} = \frac{C_{master}^{(I)}}{l_T(\alpha)} \quad (3)$$

Here, the superscripts (II) and (I), respectively, indicate the next (second) generation, and previous (first) generation. By adopting a next-generation master router with capacity indicated by (3), we maintain the same relative master-slave scaling, and hence the same optimum energy-efficiency gain. In the following analysis, we calculate two possible upgrade trajectory scenarios, based on α=0.1 and α=0.3. Fig. 3(a) shows the upgrade path for α=0.1 as we increase overall edge-router node capacity from 1.05 Gb/s through to 1 Tb/s. Our baseline assumes a 1-Gb/s router consumes of the order of 100 W power. Each generation is shown as a different colour, where each curve is the appropriate plot of Eqn.(1), representing the expected average power consumption as l_T is varied for the particular master-slave configuration. The minimum in each curve (e.g. for generation (I) this is at 32.9 W @ 190 Mb/s) represents the optimum operating point for l_T; and hence in this case the average (expected) power consumption for a configuration with maximum master capacity at the high end of the curve (1.05 Gb/s) and slave capacity at that minimum point (190 Mb/s), i.e. $l_T = 0.18$. The high end of each curve coincides with the minimum (i.e. optimum operating point) of the next-generation curve. Hence the master capacity of the previous generation becomes the slave capacity of the subsequent upgrade. In such a way we follow the $P \propto C^{2/3}$ relationship (i.e. the main diagonal of the graph) but continue to enjoy an average 68.1% energy-efficiency gain. Five upgrade generations are required to attain a 1 Tb/s ARN edge-router capacity, e.g, as may be required in the 2030 timeframe. Fig.3(b) shows a different upgrade trajectory assuming a higher average traffic intensity of α=0.3, consisting of nine generations to upgrade from 655 Mb/s to 1 Tb/s, but with a lower average energy-efficiency gain of now only 45.7%.

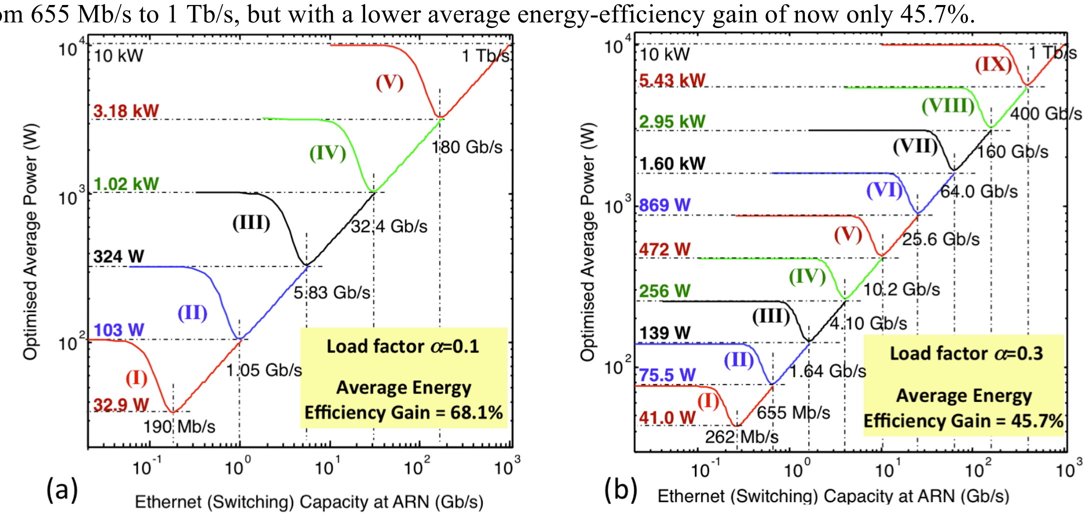

Fig.3. Upgrade trajectories from 1Gb/s→1Tb/s: (a) 5 generations for α=0.1, efficiency gain η=68.1%; (b) 9 generations for α=0.3, and η=45.7%.

4. Discussion

Two differing upgrade trajectories have been presented here, offering very different migration profiles. While Fig. 3(a) offers fewer upgrade stages and a higher energy-efficiency gain, the system is highly over-provisioned; whereas Fig. 3(b) indicates an approach that runs closer to maximum capacity for each generation, but requires more frequent upgrades and offers a lower energy-efficiency gain. A fuller techno-economic comparative analysis of these two migration paths is the subject of another paper.

5. Acknowledgement

This work was supported by the EU OASE project no. 249025, and the EU STRONGEST project no. 247674.

6. References

[1] W. Van Heddeghem, F. Idzikowski, W.Vereecken, D. Colle, M. Pickavet and P. Demeester, "Power consumption modeling in optical multilayer networks", Photonic Network Communications, January 2012 (2012)

[2] P.E. Ross, "5 Commandments [technology laws and rules of thumb]", IEEE Spectrum, **12**, 30 (2003).

[3] R.S. Tucker, "Optical packet-switched WDM networks: a cost and energy perspective", OFC'08, Paper OMG1, San Diego (2008).

[4] S. Aleksić, "Power consumption of hybrid optical switches", Proc. OFC'10, Paper OThP6, San Diego (2010).

[5] M.C. Parker, S.D. Walker, "Stochastic Energy-Efficiency Optimization in Photonic Networking by use of Master-Slave Equipment Configurations", *ONDM'12*, Paper 7Wed4, Colchester, UK, (2012)

[6] M.C. Parker, S.D. Walker, 'Roadmapping ICT: An Absolute Energy Efficiency Metric ', JOCN, **3(8)**, A49 (2011)

[7] G. Shen, R.S. Tucker, "Energy-minimized design for IP over WDM networks", JOCN, **1(1)**, p176 (2009).

[8] W. Van Heddeghem, F. Idzikowski, E. Le Rouzic, J. Y. Mazeas, H. Poignant, S. Salaun, B. Lannoo and D. Colle, "Evaluation of Power Rating of Core Network Equipment in Practical Deployments", submitted to IEEE greencom 2012, online conference, (2012)