Burst Switching for Energy Efficiency in Optical Networks

Shuping Peng, Kerry Hinton*, Jayant Baliga*[+], Rodney S. Tucker*, Zhengbin Li and Anshi Xu
State Key Laboratory of Advanced Optical Communication Systems & Networks (Peking University)
Beijing 100871, P. R. China
** Centre for Ultra-Broadband Information Networks (CUBIN), [+]National ICT Australia*
University of Melbourne, Vic 3010, Australia
pengs@unimelb.edu.au

Abstract: Energy consumption of electronic burst switching is modeled and compared to electronic packet switching in the network core and edge. It is shown that burst switching can provide significant energy savings relative to packet switching.

OCIS codes: (060.4250) Networks; (220.4830) Systems design

1. Introduction

As traffic increases, the possibility emerges that Internet growth may ultimately be constrained by energy consumption rather than by bandwidth [1, 2]. Network operators are paying increased attention to power consumption, and the power consumption of future high-capacity packet switching routers is a growing problem [3]. The power consumption in a packet switching router is dominated by the IP header processing and forwarding [4]. A potential alternative to packet switching is burst switching (BS) [5]. BS significantly reduces the amount of header processing in the switch and so may provide an avenue for addressing the power consumption issue. To test this, we construct an energy consumption model of an electronic burst switch based on commercially available sub-systems. Using this model to compare power consumption of burst switches with electronic packet switches, we find that a core burst switch could consume less than half of the power consumed by a core packet switch, and an edge burst switch would consume little more power than an edge packet switch.

2. Network Architecture

To estimate the power consumption in a burst and packet switches, we break down their functions into sub-systems for which we can realistically estimate power consumption values. We distinguish between edge and core switches. The *BS edge switches* aggregate packets into bursts and generate the burst control packets (BCPs). The *BS core switches* route bursts through the core network. To construct power consumption models, we present typical core and edge switch architectures for electronic burst and electronic packet switches [4].

2.1 Core switch architecture

A core BS architecture is shown in Fig. 1. There are three main functional blocks: a) the *switch control unit* (SCU), b) the *line card* and c) the *switch fabric*.

The SCU structure is shown in Fig. 3. BCPs processing and routing functions are undertaken in the SCU, so the BS line cards are simpler than in packet routers. Moreover, less processing occurs in the SCU, since generally hundreds of packets (10 kbits per packet) are assembled into each burst (1 Mbits per burst) by edge burst switch, and only BCPs are processed in a core BS. The SCU also includes a routing & forwarding (R & F) table (not shown in the figure). The R & F table contains the routing information and the current configuration status of the switch fabric. The information in the BCP (i.e. source/destination address, burst length, offset time and QoS) is extracted and processed. The burst is switched as instructed by this information. Finally, the information in the BCP is modified and it is transmitted to the next node. We can utilize high capacity commercial single chip and high speed electronic interconnects to build a high capacity switch fabric. For scalability, a 3-stage CLOS architecture [4] can be adopted.

A typical packet switch architecture is shown in Fig. 2. The main differences between the burst and packet switch architectures are the processes undertaken in the line cards and SCU. The SCU in the packet switch contains a routing engine, which maintains the routing table. In the packet switch line card, shown in Fig. 4, there is a forwarding engine, which manages data flow to the switch fabric and the buffers. After the incoming packet is received, the information in the packet header is extracted and processed, while the payload is buffered. The packet is segmented into cells which are routed through the switch fabric as instructed by the routing engine. In the output port, the cells are reconstructed into packets, the new header is attached and the packets are sent to the output interfaces.

2.2 Edge switch architecture

In an edge burst switch, incoming packets are sorted according to destination and CoS, and assembled into bursts. When bursts are assembled, BCPs are generated and sent to core nodes. At the same time, an offset_time timer is started up. After the offset time, bursts are sent to core nodes along the route reserved by BCPs. There are more buffers in the edge burst switch than in an edge packet switch because they are required for bursts (dis)assembly. Burst (dis)assembly means the processing in the edge burst switch is also more complicated than in an edge packet switch.

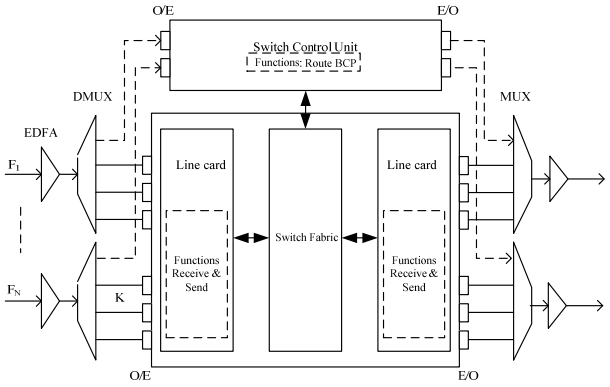

Fig. 1. Core node architecture for Burst Switching

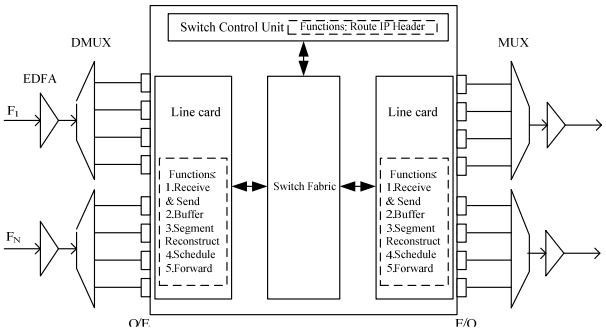

Fig. 2. Core node architecture for Packet Switching

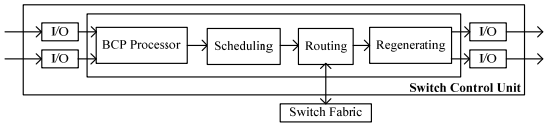

Fig. 3. Switch Control Unit for Burst Switching

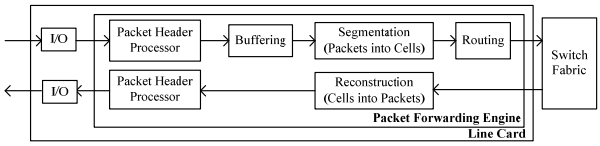

Fig. 4. Line Card for Packet Switching

3. Energy per bit Comparison

We base our energy consumption estimates on the Cisco CRS-1 multi-shelf router. This packet router has a switching capacity 92 Tb/s and line rate 40 Gb/s with a total power consumption is 1020 kW [6]. Breakdown of power consumption of a CRS-1 is shown in Table I.

Table I. Breakdown of Power consumed by Cisco CRS-1 Multi-shelf Router

		Percentage of Total Power [7]	Energy per bit
Line card	I/O Interfaces	6.5	0.7 nJ/b
	Forwarding Engine	32	3.5 nJ/b
	Buffers	3.5	0.4 nJ/b
Switch Fabric		14.5	1.6 nJ/b
Switch Control	Routing Engine	10.5	1.1 nJ/b
Power Supply & Blowers		33	3.6 nJ/b
Total		100	10.9 nJ/b

The energy per bit comparison between the two core switches is shown in Table II. In line cards of the core burst switch, there are only I/O interfaces which, from Table I, consume 0.7 nJ/b. Packet switch line cards include I/O interfaces, the forwarding engine and buffers which consume 0.7 nJ/b, 3.5 nJ/b and 0.4 nJ/b, respectively. Switch fabrics of the two switches are the same, and consume 1.6 nJ/b. In the core burst switch, the SCU performs similar functions as the forwarding engine in the core packet switch. Therefore, we assume the core burst switch SCU also consumes 3.5 nJ/b. Further, the core burst switch SCU is separate to the switch fabric, so there are additional I/O interfaces in the SCU of a core burst switch. The SCU of a packet switch contains a routing engine, which consumes 1.1 nJ/b. The power supply and blowers consume 33% of the total power consumed by the switches [7].

Because bursts are assembled from hundreds of packets and BCP processing occurs only once per burst, the energy consumption per bit of the SCU in the BS is significantly reduced. Therefore, the total energy per bit of the core burst switch is shown separately for the SCU (4.2 nJ/b) and the rest of the burst switch (5.5 nJ/b). The total energy per bit for a core packet router is 10.9 nJ/b.

The energy per bit comparison of the edge switches is shown in Table III. There are several stages of packet processing in an edge switch. All the incoming packets are initially processed at the IP level. This is in the *IP Packet Processing Stage*. So routing engine power consumption of 1.1 nJ/b is assumed for both switches. In the edge burst switch, the processed packets are sorted to different queues according to destination and CoS in the *Sorting Stage*. We adopt two sets of buffers, one for destination sorting and the other for CoS sorting. In the edge burst switch, there is an extra *(Dis)Assembling Stage*, in which the sorted packets are assembled or disassembled into corresponding buffers. In the packet switch, there are only input and output buffers. We can see that there are at least four times as many buffers (each consuming 0.4nJ/b) in edge burst switch than in edge packet switch. A switch fabric is assumed in the *Scheduling Stage* for scheduling bursts to the appropriate output ports. We can also include output buffers in the edge burst switch in case that there is no available channel when bursts are sent out. Finally, we can get the total energy per bit: 8.5 nJ/b for the edge burst switch and 6.7 nJ/b for the edge packet switch.

OWY5.pdf

Table II. Energy per bit comparison between core switches

	Burst Switch		Packet Switch	
Line card	I/O	0.7nJ/b	I/O	0.7nJ/b
	--	--	Forwarding Engine	3.5nJ/b
	--	--	Buffers	0.4nJ/b
Switch Fabric		1.6nJ/b		1.6nJ/b
SCU	Forwarding Engine	3.5nJ/b	Routing Engine	1.1nJ/b
	I/O	0.7nJ/b	--	--
Power Supply & Blowers		3.2nJ/b		3.6nJ/b
Total Energy per bit	SCU	4.2nJ/b		10.9nJ/b
	Rest	5.5nJ/b		

Table III. Energy per bit comparison between edge switches

	Burst Switch		Packet Switch	
IP Packet Processing	I/O	0.7nJ/b	I/O	0.7nJ/b
	Routing Engine	1.1nJ/b	Routing Engine	1.1nJ/b
Sorting	Buffers (Desti.)	0.4nJ/b	Buffers	0.4nJ/b
	Buffers (CoS)	0.4nJ/b	--	--
(Dis)Assembling	Buffers (Assem.)	0.4nJ/b	--	--
	Buffers (Disassem.)	0.4nJ/b	--	--
Scheduling	Switch Fabric	1.6nJ/b	Switch Fabric	1.6nJ/b
Buffering (Output)	Buffers (Optional)	0.4nJ/b	--	--
	I/O	0.7nJ/b	I/O	0.7nJ/b
Power Supply & Blowers		2.8nJ/b		2.2nJ/b
Total Energy per bit		8.5nJ/b		6.7nJ/b

4. Power Consumption Comparison

The total power consumption of core burst and packet switches for throughputs up to 100 Tb/s is shown in Fig. 5 (a) and for the edge switches in Fig. 5 (b). We have assumed that, on average, each burst consists of 100 packets. Fig. 5 (a) shows that the core burst switch consumes approximately half of the power consumed by the packet switch for a given throughput. This is because the burst switching undertakes significantly less header processing than the packet switching. Fig. 5 (b) shows that the edge burst switch consumes approximately 20% more power than the edge. This is due to the burst assembly processes in the edge burst switch.

We also investigate the impact of different burst lengths on the total power consumption of the core burst switch, which is shown in Fig. 5(c). Each burst consists of 10 (100 kbits), 100 (1 Mbits) and 1000 (10 Mbits) packets, respectively. Results show that as the burst length increases, the power consumption of the burst switch is reduced. However, when the burst length is larger than 1 Mbits, the power consumption is only marginally reduced.

From these results, we find that provided the number of edge switches is less than 4 times the number of core switches, then burst switching will provide a lower total network consumption than packet switching. In this, we assume the energy consumption of optical transport between the switching nodes is identical in the two networks.

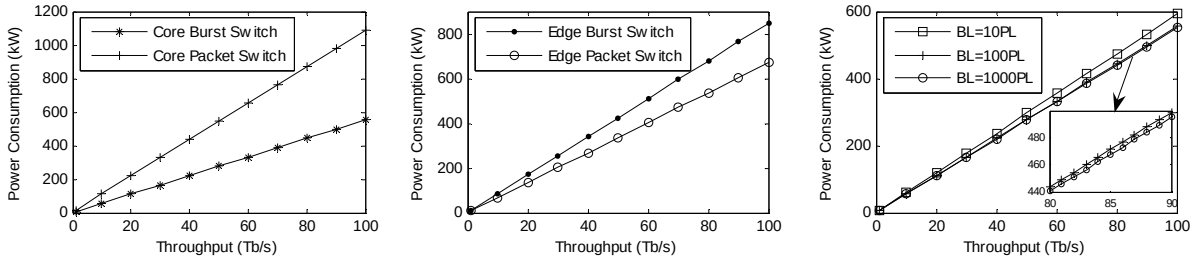

(a) (b) (c) BL: Burst Length, PL: Packet Length

Fig. 5. Power consumption of the burst switching and the packet switching (a) Core switch (b) Edge switch (c) Different burst lengths

5. Conclusions

Using power consumption data for commercially available equipment, we have compared the power consumption of electronic burst and packet switch architectures. Results indicate that significant power savings can be attained in core burst switch relative to the core packet switch. However, the edge burst switch consumes approximately 20% more power than an edge packet switch. These results indicate that burst switching networks can attain improved energy efficiency compared to packet switching networks.

Acknowledgements

This work is supported by National Basic Research Program of China (973) No. 2010CB328201, 2010CB328202 and 2010CB328203; National High Technology Research & Development Program of China (863) No. 2008AA01A329; NSF of China No. 60372025; International Academic Exchange Fund of Graduate School of Peking University.

References

[1] M. Gupta and S. Singh, "Greening of the Internet", *ACM SIGCOMM*, Karlsruhe, Germany, Aug. 2003.

[2] J. Baliga, R. Ayre, K. Hinton, W.V. Sorin, R.S. Tucker, "Energy Consumption in Optical IP Networks", *IEEE/OSA J. Lightwave Technology*, vol. 27, Issue. 13, pp. 2391-2403, 2009.

[3] D. Neilson, "Photonics for Switching and Routing", *IEEE JSTQE*, vol. 12, no. 4, pp. 669-678, 2006.

[4] R.S. Tucker, R. Parthiban, J. Baliga, K. Hinton, R. Ayre, W.V. Sorin, "Evolution of WDM Optical IP Networks: A Cost and Energy Perspective", *IEEE/OSA J. Lightwave Technology*, vol. 27, Issue. 3, pp. 243-252, 2009.

[5] S. Amstutz, "Burst Switching – An Introduction", *IEEE Communications Magazine*, pp. 50-57, Sept.1989.

[6] Cisco Data Sheets, http://www.cisco.com.

[7] J. Baliga, R. Ayre, K. Hinton, R.S. Tucker, "Photonic Switching and the Energy Bottleneck", *Photonics in Switching*, pp. 125-126, 2007.