

Transparent, Low Power Optical WDM Interface for Off-Chip Interconnects

Odile Liboiron-Ladouceur, Howard Wang, and Keren Bergman
Columbia University, Department of Electrical Engineering
500 West 120th Street, New York, NY 10027 USA

Abstract- We demonstrate a scalable WDM optical interface with 70% improvement in static power efficiency for off-chip interconnects. Serial 40Bytes packet streams are transparently mapped onto a WDM optical signal with a measured power penalty of 1.5 dB.

I. INTRODUCTION

The recent emergence of chip multiprocessors (CMPs) driving performance enhancement via increases in the number of parallel computational cores has accelerated the bandwidth requirements in high-performance processors [1]. With the vastly growing number of cores and increasing on-chip computation, communications on-chip and most critically off-chip has become the key bottleneck limiting system performance. A high-bandwidth, low-latency, and, perhaps most importantly, low-power communication infrastructure is clearly required for next generation high-performance computing systems.

Optical interconnects offer a potentially disruptive technology solution for directly addressing the bandwidth and power limitations of electronic interconnects [2]. Parallel optical links for board-level inter-chip optical communication or inter-board communication through a backplane have been recently demonstrated to offer impressively high data throughput with relatively low power consumption [3]. However, as off-chip bandwidth demands of CMPs continue to accelerate (a current generation Cell processor requires 50.6GBytes/s [4]), the power dissipation associated with multiple parallel electro/optic signal conversions grows rapidly.

In this work, we propose and demonstrate a scalable and transparent interface approach that directly maps serial streams of electronic packets onto multiple WDM channels in a highly power-efficient manner. The optical interface exhibits significant power savings since only one high-speed optical modulator, one broadband gate and one optical receiver are required for a number of WDM channels. In the experiments reported here we show a 70% improvement in the static power dissipation of the interface. The demonstrated interface translates serial 40Byte PCI Express (PCIe) encoded electronic packets onto an 8-channel WDM optical link with a measured power penalty of 1.5 dB.

The WDM optical signals are envisioned to propagate through embedded waveguides within a backplane or over single-mode fiber links between racks (Fig. 1). The PCIe protocol is employed in this experiment as it has

clearly emerged as the leading I/O interconnection standard supporting high-performance chip-to-chip and board-to-board applications [5]. PCIe was further chosen to demonstrate the flexibility and scalability (to 32 lanes) of the interface.




Fig. 1 Transparent WDM interface for inter-node communication infrastructure.

II. Reducing Power Consumption

In parallel optical links [3], each optical channel requires its own set of drivers, optical modulators, photodetectors, transimpedance and limiting amplifiers, and other related circuitry (e.g. for thermal stability). In CMOS technology [6], the total power dissipated ($P_{tot,elec}$) by the components for the signal conversion corresponds to the number of WDM channels (N) multiplied by the sum of the static power and the dynamic power as expressed in equation 1.

$$P_{tot,elec} = N(P_{static} + CV_{DD}^2 f_s) \quad (1)$$

Leakage current in digital logic (e.g. drivers) and bias current in analogue circuitry (e.g. amplifiers) associated with the optical modulator and receiver contribute to the overall static power dissipation of the interface. The dynamic power is proportional to the gate capacitance (C) and the voltage supply (V_{DD}), as well as the switching frequency (f_s).

Compared with parallel optical links, the transparent WDM optical interface partitions serial packets onto multi-channel packets in a more power efficient manner. While the dynamic power remains the same, there is significant static power savings due to the use of only one modulator/receiver pair. In fact, the power savings attained via the WDM interface ($P_{tot,WDM}$) compared to parallel optical links ($P_{tot,elec}$) increases with the number of WDM channels.

$$\frac{P_{tot,WDM}}{P_{tot,elec}} = \frac{P_{static} + N \cdot CV_{DD}^2 f_s}{N \cdot P_{static} + N \cdot CV_{DD}^2 f_s} \quad (2)$$




Fig. 2: Schematic of the WDM interface for off-chip.

III. Transparent WDM Interface Design

In Fig. 2, the off-chip electronic packets originating from the processor are directly mapped onto multiple WDM channels (W0 to W7) by simultaneously modulating all wavelengths using one high-speed optical modulator (MOD). Each modulated wavelength is then filtered using 100 GHz thin film filters (TFF) and appropriately delayed with respect to each other using optical fiber delay lines (FDL) by an amount corresponding to the predetermined WDM packet segment length. All wavelengths are then multiplexed onto one fiber. A broadband optical modulator (Gate) precisely gates the optical signals for the WDM packets as shown in Fig. 3. For a multi-lane serial protocol, such as PCIe, each lane would be interleaved in the created gap. At the destination core, the serial packet is optically reconstructed by filtering and delaying each wavelength in a manner complementary to that at the originating node. The WDM channels are multiplexed onto one fiber prior to detection at a DC-coupled broadband receiver (RX) enabling the reconstruction of the original serial streams of electronic packets.




Fig. 3: Serial electronic packets onto multiple WDM channels.

In this demonstration, a serial 8b/10b PCIe packet of 40 bytes modulated at 2.5 Gbps was mapped onto 8 WDM channels spaced by 0.8 nm (1543.72 nm to 1549.33 nm), each containing 5 bytes of the serial electronic packet [7]. A power savings of 73.5% (N=8) was achieved when taking into account the broadband optical modulator. In terms of performance, the power penalty of the transparent interface was measured with respect to a back-to-back link with the entire electronic packet encoded on one WDM channel

(W0). In Fig. 4, the observed 1.5 dB power penalty (BER 10^{-12}) is a function of the system's ability to maintain equal power and extinction ratio across all the WDM channels for the one fixed voltage threshold used by the optical receiver.

The transparency of the optical interface enables mapping of various serial protocols (e.g. RapidIO). Longer serial data streams can be mapped by increasing the number of WDM channels. For example, 1024 bytes can be mapped onto 16 WDM channels each containing 512 bits modulated at 10 Gb/s with FDL length increment of 51.2 ns. With recent advances in silicon photonics [8] and slow-light techniques to delay the packets [9], we can envision the transparent WDM interface integrated within a processor die providing a power efficient, low-latency communication infrastructure for both on-chip and off-chip communications in CMPs.




Fig. 4: Measured power penalty of the interface.

IV. Conclusions

A scalable and transparent WDM optical interface was demonstrated with 70% improvement in the static power efficiency. PCIe encoded electronic packets of 40Bytes were mapped onto 8-channel WDM optical link with a power penalty of 1.5 dB. The low-power WDM interface efficiently addresses the growing demands of off-chip bandwidth CMPs.

This work was supported by the U.S. Department of Defense under subcontract B-12-664.

REFERENCES

- [1] S. Vangal *et al.*, "An 80-Tile 1.28TFLOPS Network-on-Chip in 65nm CMOS," *ISSCC Dig. Tech. Papers*, pp. 98-99, Feb. 2007.
- [2] D. Miller, "Rationale and Challenges for Optical Interconnects to Electronic Chips," *Proc. IEEE*, vol. 88, no. 6, pp. 728-749, 2000.
- [3] C. Schow *et al.*, "160-Gb/s, 16-Channel Full-Duplex, Single-Chip CMOS Optical Transceiver," *Proc. OFC*, OThG4, Mar. 2007.
- [4] M. Kistler *et al.*, "Cell Multiprocessor Communication Network: Built for Speed," *IEEE Micro*, vol. 26, no. 3, pp.10-23, May-Jun. 2006.
- [5] Budruk *et al.*, "PCI Express System Architecture," *Addison Wesley*, Sept. 2003.
- [6] ITRS- 2005 edition, <http://public.itrs.net>.
- [7] O. Liboiron-Ladouceur *et al.*, "An All-Optical PCI-Express Network Interface for Optical Packet Switched Networks," *Proc. OFC*, JWA59, Mar. 2007.
- [8] R. Soref, "The Past, Present, and Future of Silicon Photonics," *IEEE JSTQE*, vol. 12, no. 6, pp. 1678-1687, Nov.-Dec. 2006.
- [9] R.S. Tucker *et al.*, "Slow-light optical buffers: capabilities and fundamental limitations," *IEEE J. Lightw. Technol.*, vol. 23, no. 12, pp.4046-4066, Dec. 2005.