Reducing Energy per Delivered Bit in Silicon Photonic Interconnection Networks

Ke Wen, Sébastien Rumley, Keren Bergman
Department of Electrical Engineering, Columbia University
New York, NY 10027

***Abstract*—A novel interface design with an optically-implemented availability feedback protocol is proposed for silicon photonic networks. Simulation results show a reduction in *energy per delivered bit* by as much as 34% compared to prior schemes.**

I. Introduction

Optical switching enables data signals to remain in the optical domain and is a promising solution to energy-efficient cluster-scale interconnects. Silicon photonic (SiP) technologies present a platform for realizing such switches as well as many other optical network components (modulators, filters, detectors, etc.) using CMOS processes. However, even with significant strides made in improving the energy efficiency (EE) of individual devices (Table I), the actual energy consumption must be viewed from a network operation perspective. Improperly designed network interfaces and packet injection policies can easily negate even large efficiency improvements achieved at the device level.

A SiP communication system, unlike most electronic networks, is not energy proportional: a substantial amount of static power is consumed by source lasers and thermal stabilization which need to remain on for the lifetime of network operation. By comparison, dynamic power (modulation and detection) is consumed during data communication. While the static consumption must be amortized by high utilization, dynamic EE depends on the packet drop rate, which unfortunately increases with the utilization. It is therefore critical to evaluate the EE of network injection schemes against the metric of *energy per delivered bit (EPDB)*, rather than per transmitted bit.

SiP switching systems have used packet injection schemes that can be categorized into three classes: *prescheduled time-division multiplexing (TDM)*, *opportunistic*, and *hybrid*. Since traffic patterns can be highly nonuniform and unpredictable in real time, the above schemes all have their advantages and drawbacks. For example, prescheduled TDM achieves a 100% successful rate, but fails to adapt to real-time nonuniformity, leading to a lowered throughput and static EE. Opportunistic access (OA) retains high throughput by allocating bandwidth according to real-time demands, but suffers from uncoordinated injections and contentions, leading to lost

TABLE I
Energy efficiency (pJ/bit) of network components [6].

Microring modulator & driver	0.35	Dynamic
Photodetector & receiver circuit	1	Dynamic
Elec. transmission line (btwn buffer & Tx)	2	Dynamic
Laser (10% wall-plug efficiency)	0.70	Static
Thermal stabilizer (at 25 Gb/s)	0.06	Static

energy carried by dropped packets. The hybrid scheme [1] does incorporate OA into prescheduled TDM, which allows it to utilize the transmitter capacity of an unused prescheduled slot, but OA packets can be preempted by TDM packets in the network, creating another source of contention and energy inefficiency.

Building on OA's adaptability to varying traffic and low latency in the low-load region [1], this work proposes a novel interface design with an optically-implemented availability feedback protocol, which mitigates contentions and reduces the EPDB by as much as 34% compared to previous injection schemes.

II. Optical Network State Feedback

OA typically resolves contentions by retransmission. However, without knowledge of network states, such retransmission can create head-of-line (HOL) blockings. While frequent retrials (e.g. FastNACK [2]) waste dynamic energy, backoff methods lower the static EE. The key to resolving this problem is the feedback of network states. While adding an electrical control plane across the whole system is cumbersome, we present an *optical availability feedback* (OAF) scheme based on the SPINet architecture [3]. OAF enhances EE by 1) fast terminating rejected transmissions and 2) avoiding HOL blockings.

Depending on the physical limitation, a SiP interconnect can comprise a single or multiple SiP chips, with each chip constituting a switching element (SE) in a multistage network (Fig. 1a). SPINet employs a frame wavelength (λ_F, sending '1' during packet existence) and header wavelengths (λ_H, representing packet address) for fast switch configurations. OAF makes use of this feature to transmit availability information from the switch interface back to the client. As Fig. 1b shows, an *admission and feedback* module is added to each input

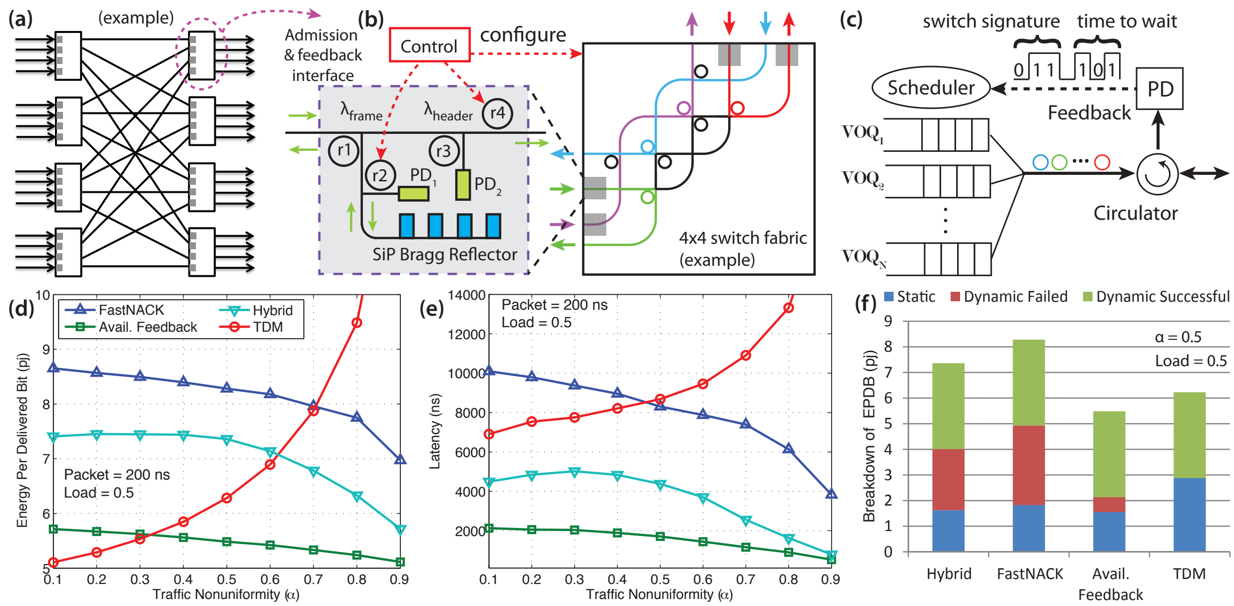

Fig. 1. (a) An example multistage interconnect formed by (b) an example 4×4 SiP switch element incorporated with the admission and feedback interface; (c) client Tx interface; (d) energy efficiency and (e) latency versus traffic nonuniformity α; (f) energy profile when $\alpha = 0.5$.

of the switch. Microring **r1** and **r3** filter out λ_F and λ_H, respectively. Microring **r2** is initially tuned to on-resonance state for the detection of λ_F. When there is no contention, the control logic configures the switch fabric and tunes **r4** to pass the input signal to the switch fabric unperturbed. Otherwise, if the packet is not accepted, **r2** will work as an active modulator driven by the control logic for OAF. When **r2** is tuned to the off-resonance state, it passes λ_F to a SiP Bragg reflector [4], which reflects the optical power back to the client as a '1'. Otherwise, **r2** drops the optical power of λ_F to detector $\mathbf{PD}_1$ and the client will see a '0'.

The feedback from a blocking SE to a client includes two parts: the SE "signature" and the time to wait (TTW) before the blocking port becomes available (Fig. 1c). The client, with a virtual output queue (VOQ) structure and knowledge about the blocking position, disables not only the rejected VOQ but also those *whose packets pass through the same position*. TTW timers of these VOQs are updated to the TTW value received. A VOQ will not be scheduled until its TTW timer expires, avoiding blockings. Although OAF requires detection of feedbacks on λ_F, its energy overhead can be easily amortized by the large number of parallel data wavelengths used.

III. Simulation Results

EPDBs of various injection schemes are evaluated by simulation on a 32-node omega network. Each client has the same offered load but with a destination distribution biased towards its nearest neighbor. Such nonuniformity can be described by the (α, S) model [5]: a source-destination stream is said to be (α, S) if it contains at most αS packets in any S consecutive slots ($0 < \alpha \leq 1$). EEs of network components are listed in Table I [6].

Since TDM (preconfigured for uniform traffic) suffers from lowered throughput as the traffic becomes nonuniform, its EPDB increases with α (Fig. 1d). In comparison, EPDBs of OA-involved schemes decrease with α. Specifically, OAF reduces the EPDB by as much as 34% and 25% compared to FastNACK and *Hybrid*, respectively. *W.r.t.* TDM, OAF consumes slightly more energy in the small-α region ($\alpha \leq 0.3$), but offers tremendous efficiency advantage over TDM thereafter. Fig. 1f shows the EPDB profiles of different schemes when $\alpha = 0.5$. It explains why OAF has a lower EPDB: OAF saves on the dynamic energy spent on failed (re)transmissions compared to other OA schemes, while it also has a lower *static EPDB* than that of TDM due to a higher throughput. Fig. 1e further shows that OAF also achieves much lower latency than all other schemes.

IV. Conclusion

The EPDB of packet injection schemes in SiP networks is related to traffic nonuniformity. The proposed OAF protocol effectively avoids HOL blockings and significantly improves the EPDB as well as latency.

References

[1] Chrysos, *IEEE HPSR*, 2010.
[2] Dongaonkar, *Opt. Interconnects*, 2013.
[3] Shacham, *IEEE Micro*, vol. 27, no. 4, 2007.
[4] Giuntoni, *Opt. Express*, vol. 17, no. 21, 2009.
[5] Weller, *IEEE/ACM Trans. Netw.*, vol. 5, no. 6, 1997.
[6] Ophir, *IEEE Micro*, vol. 33, no. 1, 2013.