Towards Chip-Scale Plasmonic Interconnects

Hassan M. G. Wassel[1], Mohit Tiwari[1], Jonathan K. Valamehr[2], Luke Theogarajan[2], Jennifer Dionne[3], Frederic T. Chong[1], and Timothy Sherwood[1]

[1]Department of Computer Science, UC Santa Barbara
[2]Department of Electrical and Computer Engineering, UC Santa Barbara
[3]Department of Materials Science and Engineering, Stanford University

1 Introduction

In the multi and many core era, communication is crucial to the system performance. Therefore, network-on-chip (NoC) approaches were proposed to regularize the design of on-chip communication. Nanophotonic interconnects have been proposed in the recent few years as a replacement of global metal interconnect because of their almost distance-independent power consumption and low-latency and high bandwidth. However, photonic components have their limitations: diffraction-limited sizes, temperature dependence and cross-talk and bend losses.

First, nanophotonic components are governed by the diffraction limit which dictates that light cannot be confined in a space smaller than $\lambda/2n$ where λ is the free space wavelength and n is the refractive index of the material. This means that the size of waveguides and modulators and ring filters are in the micrometer scale because of usage of C-band around 1550 nm wavelength. This size mismatch between micrometer-scale photonic components and nanometer-scale electronic ones limits the integration viability of both technologies on the same chip using the same process. These relatively bigger size components have higher capacitance which requires more power consumption to drive and limits the speed at which they can operate. For example, a photonic link is expected to have around 158 fJ/bit electric and electro-optical components energy consumption per bit [1] which is clearly over the estimated viability requirement of 10 fJ/bit per device [7]. Second, photonic components are temperature-dependent which is exploited in the micro-ring modulators that their resonance range is adjusted using heating. This heating requirement is estimated to consume around 100 fJ/bit [1]. These two major power consumption components limit the minimum distance at which nanophotonic waveguides can be more power efficient than electrical signaling, even excluding off-laser generator power consumption. Third, because of low confinement, bend losses and cross-talk are a problem for nanophotonic devices. The minimum pitch of a Si waveguide is 5.5 μm. This paper introduces a new technology for interconnects, plasmonics, that provides an alternative to a purely photonic interconnect for many-core processors.

2 Surface Plasmon Polaritons

Surface plasmon polaritons (SPP) are electromagnetic waves that are coupled to free electron collective oscillations in a metal. When a light beam is incident a metal-dielectric interface, surface plasmon polaritons are excited and will propagate along the surface of the metal. Interestingly, surface plasmons excited at the interface of a metal and dielectric will maintain the frequency of the exciting light, while at the same time have a much shorter wavelength. These shorter wavelengths allow the construction of nanoscale devices that tightly confine even very high frequency electromagnetic waves, in a way side-stepping traditional diffraction limits. Of course nothing comes for free, and SPP propagation is limited by both metal absorption (i.e. ohmic losses) and free-space radiation. During the last decade, the plasmonics field has made significant progress in improving the propagation distances of SPP modes and recently, we have seen a flurry of new work on both active and passive components and plasmonic sources.

3 Plasmonics Model

Plasmonic Components

A plasmon slot waveguide (or metal-dielectric-metal (MDM) waveguide) consists of two metal sheets separated by a thin dielectric core [4]. The thickness of the dielectric and the skin depth of the metal determine the wavelength of the light signal that can propagate in the waveguide. For maximum mode confinement, the depth of the metal should be optically opaque, as thick as the metal skin depth (approximately 20 nm). Having a metal thickness a little higher than double the skin depth ensures that modal fields remain confined within the core and metallic cladding, leading to very low pitch of 100s of nanometers (compared to 5.5 μm of Si waveguides). The dielectric thickness can range from just a few nanometers to over hundreds of nanometers. For a core thickness less than roughly 50 nm, only plasmonic modes will propagate. As the core thickness is increased, the waveguide begins to support transverse electric and magnetic photonic modes, similar to microwave waveguides.

In a recent study of Si-based plasmonics [3], it is found that in terms of propagation distances, air is the best dielectric followed by SiO_2 and then Si. For $\lambda = 1550$ nm, signals can propagate up to 80 μm in a silver/silicon oxide/silver waveguide with 250 nm core [4]. It is worth noting that an active area of plasmonics research is the introduction of gain material in the dielectric core of the slot waveguide to increase the propagation distance. For the sake of our power analysis, we assume that a silver/silicon oxide/silver with core thickness of 10 nm has loss of 0.2 dB/μm and group velocity of c/3.75 at 1550 nm [3].

An MDM compact plasmonic modulator based on resonance in a nano slit-groove that is created in the metal layer and filled with dielectric material is extremely energy-efficient with 1 fJ/bit energy consumption [2]. Modulation is achieved by modulating the refractive index of the dielectric in the groove. The active media in the cavity can be quantum structures Stark Effect (QSSE) that was realized in a silicon-germanium system. That makes this plasmonic modulator a CMOS compatible one.

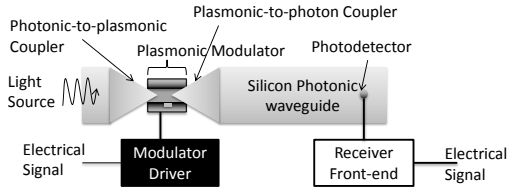

Figure 1: Hybrid Link: By using a plasmonic modulator and silicon photonic waveguide, we can achieve the best of both worlds: long range propagation, and low power consumption and high performance modulation. Couplers converts photons into SPP and vice versa.

Electrical Power Estimation

We used community standard Orion 2.0 [5] to estimate the power consumption of electrical links. We used 32 nm at Vdd = 1 V high performance transistors at 3 GHz. Orion reports power consumption that we divide over the clock rate to get the energy per bit power consumption. We assumed activity factor of 1 in all technologies.

Photonic Link Power Estimation

Transmitters are ring resonator modulators driven by an analog circuit and receivers are a Ge photodetector connected to a TIA and an amplifier. Using parameters of analog components at 22 nm technology from [1], we estimate that analog components will consume 88 fJ/bit. The modulator is estimated to consume 82 fJ/bit at 32 Gbps as in [6] and 100 fJ/bit heating power as in [1]. Detector sensitivity is assumed to be 10 μW.

Plasmonic Link Power Estimation

Because of the limited propagation distance of plasmonic waveguides to around 100 μm, we decided to connect more than one plasmonic link in order to achieve higher propagation distances. This means that we will have a modulator and a detector for each link, with detector of the current link driving the modulator of the next link. Laser power is provided by a single photonic Si-waveguide for each sub-link and using a coupler between the plasmonic and photonic waveguides. We assume 80 fJ/bit energy consumption of electric components driving each sub-link. We conservatively assume 2 dB coupler loss although higher coupling efficiencies have been demonstrated. We can calculate laser power consumption for each sub link using 10 μW detector sensitivity.

Hybrid Link Power Estimation

We propose exploiting active plasmonic modulation to modulate light in a photonic waveguide as suggested in [2] by coupling the the Si conventional waveguide to the modulator and use the modulated plasmon to couple it back into photons as shown in Figure 1. This approach saves a lot of heating power required for photonic modulators that also consumes at least one order of magnitude higher energy per bit. It uses the same electric components as the photonic links. Laser power is estimated using losses from coupling and propagation.

Link Characterization

Using these models, we estimated the energy per bit of electrical, photonic, plasmonic and hybrid links by calculating the power consumption of all of them and dividing that by the corresponding bandwidth. We assume bandwidth of 10 Gbps for plasmonic, hybrid, and photonic links and 3 Gbps for electrical links. Figure 2 shows the energy per bit consumption against the length of the link for all four configurations. Plasmonic links are modeled using a 1 fJ/bit modulator driver and the parameters given in [1]. A plasmonic link cannot be more energy efficient than electrical and photonic links at any distances because it suffers from high laser power due to its inherent losses and distance dependent power consumption.

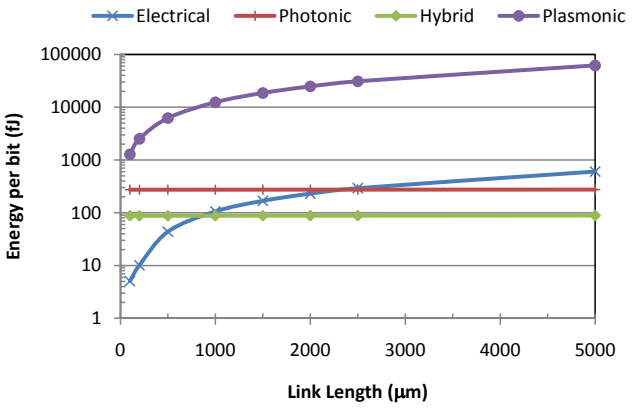

Figure 2: Link length vs Energy per bit for different technologies: electrical, photonic, plasmonic and hybrid. It is clear that electrical signaling is more efficient for any link of less than 800 μm length. Beyond that, hybrid links are the most energy efficient.

Using the plasmonic modulator with the conventional waveguide (hybrid) reduces the distance at which photonics become more energy efficient than electrical from 2.5 mm to less than 0.8 mm. This is an interesting result because we will use the best of both worlds, modulating using a small efficient plasmonic device and propagating the signal in a low-loss medium. It is worth noting that we lose the ability to support wavelength-division multiplexing because of the introduction of the plasmonic modulator that lies in the signal path whether it is modulating or not unlike ring-resonators that do not block the light if it is off-resonance.

4 Conclusions

We discuss the basics of plasmonics, and show that while pure plasmonics cannot provide a competitive energy-efficient link, a hybrid plasmonic/photonic link provides the opportunity to replace short electrical links leading to latency and energy savings.

References

[1] C. Batten, A. Joshi, J. Orcutt, A. Khilo, B. Moss, C. W. Holzwarth, M. A. Popovic, H. Li, H. I. Smith, J. L. Hoyt, F. X. Kartner, R. J. Ram, V. Stojanovic, and K. Asanovic. Building many-core processor-to-dram networks with monolithic cmos silicon photonics. *IEEE Micro*, 29(4):8–21, 2009.

[2] W. Cai, J. White, and M. Brongersma. Compact, high-speed and power-efficient electrooptic plasmonic modulators. *Nano Letters*, 9(12):4403–4411, 2009.

[3] J. Dionne, L. Sweatlock, M. Sheldon, A. Alivisatos, and H. Atwater. Silicon-based plasmonics for on-chip photonics. *Selected Topics in Quantum Electronics, IEEE Journal of*, 16(1):295 –306, jan.-feb. 2010.

[4] J. A. Dionne, L. A. Sweatlock, H. A. Atwater, and A. Polman. Plasmon slot waveguides: Towards chip-scale propagation with subwavelength-scale localization. *Physical Review B (Condensed Matter and Materials Physics)*, 73(3):035407+, 2006.

[5] A. B. Kahng, B. Li, L.-S. Peh, and K. Samadi. Orion 2.0: A fast and accurate noc power and area model for early-stage design space exploration. In *Design, Automation & Test in Europe Conference & Exhibition, 2009. DATE '09.*, pages 423–428, April 2009.

[6] N. Kirman and J. F. Martínez. A power-efficient all-optical on-chip interconnect using wavelength-based oblivious routing. In *ASPLOS '10: Proceedings of the fifteenth edition of ASPLOS on Architectural support for programming languages and operating systems*, pages 15–28, New York, NY, USA, 2010. ACM.

[7] D. Miller. Device requirements for optical interconnects to silicon chips. *Proceedings of the IEEE*, 97(7):1166 –1185, July 2009.