

Automatic Place-and-Route of emerging LED-driven wires within a monolithically-integrated CMOS+III-V process

Tushar Krishna^{§‡} Arya Balachandran^{*‡} Siau Ben Chiah^{*‡} Li Zhang[‡] Bing Wang[‡]
 Cong Wang^{*‡} Kenneth Lee Eng Kian[‡] Jurgen Michel[‡] Li-Shiuan Peh^{‡‡}

[‡]SMART LEES¹, Singapore [§]Georgia Institute of Technology, USA ^{*}NTU, Singapore ^{‡‡}NUS, Singapore

Abstract—

We leverage a recently demonstrated CMOS compatible III-V and Si monolithic integrated process to design photonic links comprising LEDs and photodiodes, as direct replacements for on-chip electrical wires. To enable VLSI-scale design of chips with such LED links, we create a library of opto-electronic standard cells, and model waveguides as traditional metal layers. This lets us integrate LED links into a commercial place-and-route tool, which treats them as electrical cells and wires for the most part, reducing design effort. We also add support for automated replacement of electrical nets with LED links.

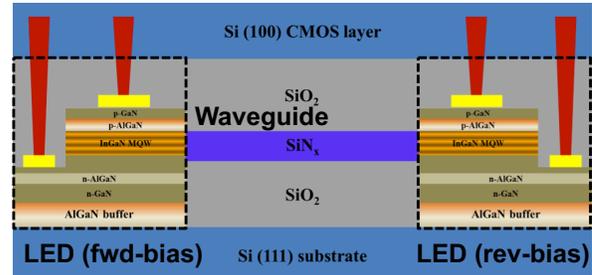
We find that LED-interconnect based designs substantially lower energy consumption vs. electrical copper wires ($\sim 39\%$ reduction in the Network-on-Chip, $\sim 27\%$ reduction within a processor core) while achieving the same latency and bandwidth, demonstrating the promise of LED on-chip interconnects.

I. INTRODUCTION

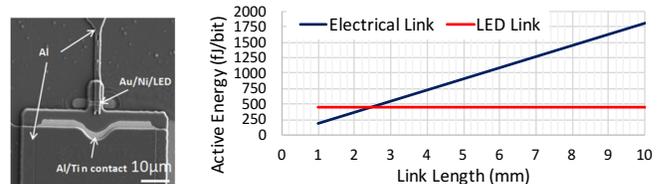
The diminishing returns from scaling of on-chip electrical interconnects compared to logic continues to remain a challenge for chip designers. Smaller transistors are faster and more power efficient, but the same does not hold true for electrical wires, making them relatively slower and more power hungry each generation. Moreover, with the advent of multicore processors and continued technology scaling (i.e., smaller tiles), the average distance of communication between on-chip cores and caches goes up each generation. Studies have shown that the energy consumed in transporting data on-chip is $10\times$ more than the energy consumed in the actual computation [1]. Electrical copper interconnects have dominated on-chip communications in commercial processors so far, and have a fundamental trade-off between energy and the interconnect length or bandwidth. New interconnects are needed to ensure scalability of future many-core processors.

Among recent disruptive technologies, optical interconnects have the potential to break the bandwidth-distance-power tradeoff of electrical interconnects [17], [22], [23]. In general, a complete optical link is comprised of a light source for generating the information carrier, a modulator for Electrical/Optical (E/O) data transformation, a photodiode for light detection, passive components for light guiding, and peripheral electronic devices for driving and biasing the photonic devices. In a photonic link, the light source is the most critical device as it consumes a substantial fraction of total link power. As Si fares poorly as a material for light emitting devices, previously proposed optical links rely predominantly on off-chip lasers as the light source. This leads to high coupling losses bringing the light on-die ($\sim 3\text{-}6\text{dB}$ [23]), which in turn

¹This work was funded by the National Research Foundation under the SMART Low Energy Electronics Systems (LEES) program.



(a) μ -LED link in LEES process [10], [15], [25]. Optical interconnects are created in a layer underneath the Si CMOS circuits.



(b) $2\times 5\mu\text{m}^2$ InGaP LED [25]. (c) Energy consumption as a function of distance in LED links and repeated electrical wires at 1Gbps.

Fig. 1: LED Interconnects in monolithic CMOS + III-V Process.

results in the need for very high optical power ($\sim \text{mW}$ [20]) that percolates downstream to modulators, photo-detectors and electrical receivers that dissipate high energy. As a result, laser-based optical links are more competitive against electrical *off-die* interconnects such as processor-to-memory IOs [23].

Recently, an alternate optical interconnect technology has been proposed and successfully demonstrated: **directly modulated on-chip μ -LEDs** in a novel CMOS compatible III-V and Si monolithic integrated process called LEES² [10], [25]. The CMOS devices are fabricated (front-end) on conventional SOI, while GaN/InGaP epitaxy is performed separately on 200 mm Si wafers before being bonded to the CMOS layers (Fig. 1). Such *on-die* μ -LEDs have potential applications as on-die interconnects due to their micro-size, more efficient usage of injected current, and much lower coupling loss directly onto an on-die waveguide ($< 1\text{dB}$ loss translating to $\sim \mu\text{W}$ optical power [16]). The process has presently been demonstrated pairing a 200 mm $0.18\ \mu\text{m}$ RF-CMOS foundry process with a $0.25\ \mu\text{m}$ critical dimension III-V process³.

In this work, we use μ -LEDs from the LEES PDK [6] to

²<http://www.circuit-innovation.org>

³The $0.18\ \mu\text{m}$ CMOS process was offered by a leading commercial foundry with capacity at that node, while the III-V process was constrained by the availability of 200 mm lithography and microfabrication resources for III-V processing. As this is a wafer-level integration process, however, it can be extended to advanced CMOS nodes on larger size wafers (e.g. 300 mm).

design optical links, and characterize the energy, performance, and area characteristics compared to electrical links. Fig. 1(c) plots the energy consumption of our designed 1Gbps μ -LED links against electrical links from post-layout simulation. All links take a single-clock cycle (1GHz) to traverse the distance specified on the x-axis. The energy consumption of μ -LED links remains the same irrespective of distance⁴, magnifying energy savings compared to electrical links as wire lengths increase. The crossover point is at 2.5mm. This is with LEDs that were over-designed for robust performance [25], and in a trailing edge CMOS process with high V_{dd}. The crossover point is expected to go down much further by optimizing material growth and device design for higher currents, designing simpler and lower power CMOS Tx/Rx circuitry, and moving to advanced nodes. Thus there is a potential for large energy benefits by replacing core-to-core and even within-core links with LED links, at no performance penalty.

There are two key challenges in realizing this goal:

- (1) Design Complexity. Integrating optical devices and E-to-O/O-to-E circuits is highly inhibitive in terms of custom layouts for each link.
- (2) Area Trade-off. Opto-electronic devices consume higher area than electrical link drivers/receivers, and it is not clear if/when the trade-off is acceptable.

This work addresses both, by the following contributions:

- We design a variety of CMOS drivers and receivers for the LED devices at various data rates, optimized for energy and area. We integrate them into an *opto-electronic standard cell library* within the LEES PDK and model waveguides as regular metal layers, hiding opto-electronic design complexity from the chip designer.
- We integrate these cells and layers into a commercial synthesis and place-and-route tool flow, laying them out like electrical cells and nets, with additional constraints related to bends and crossings to avoid optical losses.
- We create a tool to automatically identify and replace as many nets as possible in the design with LED links, subject to area constraints and the available waveguide layers.
- We perform case studies to validate the use of this toolchain both for global links across processor cores and local links within a core, and observe an energy reduction of 39% in the network and 27% within a core respectively, compared to a design with purely electrical links.

To the best of our knowledge, this is the first work to present a synthesis-to-fabrication flow for on-die LED interconnects.

Section II describes our standard-cell library of CMOS (electrical) + III-V (optical) devices. Section III presents our CAD tool. Section IV shows evaluations. Section V contrasts related work, and Section VI concludes.

II. OPTO-ELECTRONIC STANDARD CELLS

In this section, we describe the building blocks of our design (optical devices, waveguides, and electrical circuits) and how

⁴ μ -LED links consume energy for E-to-O and O-to-E conversions. The actual transmission is optical and consumes no electrical energy.

they are combined to build a standard-cell library.

A. Optical Device Design & Characterization

Light-Emitting Diodes (LED). The LEES process demonstrates that micron-size III-V LEDs on Si substrate have good light emission efficiencies which can reach $> 30\%$ external quantum efficiency. The emission intensity depends on the size of the device and injection current. The 3 dB bandwidth is around 100 MHz when operating in direct on-off modulation mode. Higher bandwidth (> 400 MHz) can be achieved by applying higher bias voltage [13].

Photo Diodes (PD). The same LED device can be operated in a reverse-biased mode as a *Photo Diode (PD)* for light detection. This eases processing and fabrication. An absorption coefficient of $10^4/\text{cm}$ [19] enables adequate O to E conversion within $\sim 10 \mu\text{m}$ absorption length. The reverse-biased current in the PD is in the nA range, making the CMOS receiver design (Section II-C) crucial for detecting the signal accurately.

GaN vs. InGaP. GaN and InGaP are two flavors of optical devices provided in the LEES PDK. We find that GaN LEDs possess shorter wavelengths and allow for tighter waveguide pitch/higher density interconnects while InGaP LEDs have a higher 3 dB modulation bandwidth (due to the faster carrier recombination rate) and lower turn-on voltage (of 2.5V). This in-turn helps lower the size (area) and energy of our CMOS Tx and Rx circuits (Section II-C) driving the InGaP LEDs. We also find that InGaP LEDs have a lower area footprint (e.g., $2 \times 5 \mu\text{m}^2$ in Fig. 1(b)) which is comparable to the area of minimum sized standard cells in the $0.18 \mu\text{m}$ CMOS process.

We thus choose InGaP LEDs as they satisfy the speed, area, and power requirements for on-chip interconnects. The PDK provides two kinds of InGaP devices that we leverage: **conservative** (high turn-on voltage, low current, and high area - targeted towards robust functional performance validated via measurement [25]), and **aggressive** (low turn-on voltage, high current, and low-area - optimized for on-chip density and currently under development).

B. Waveguide Design and Modeling

For light transmission, SiN has been demonstrated to be a suitable material having low transmission loss of ~ 1 dB/cm for the InGaP LED wavelength of ~ 670 nm [11]. The refractive index contrast between SiN and SiO₂ facilitates a compact waveguide design which is essential for dense optical interconnects. The light coupling efficiency between the light source, waveguide, and PD is good overall due to the high mode coupling efficiency of $\sim 80\%$ for the SiN/III-V transitions.

Apart from power consumption (in)sensitivity with link length, there are several key differences between optical and electrical interconnects that influence how they are best used in circuits. First, optical waveguides cannot be made finer than a certain width that depends on the wavelength of light used, as this would dramatically increase the leakage of the optical mode out of the waveguide and increase the signal propagation loss. Additionally, while optical links are electromagnetic-interference (EMI)-immune in the traditional

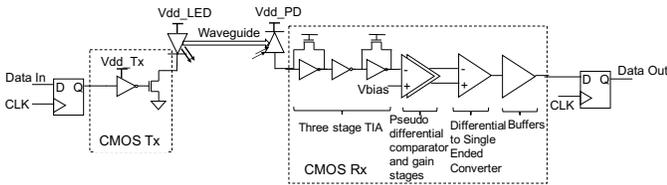


Fig. 2: Schematic of the Opto-Electronic LED Link.

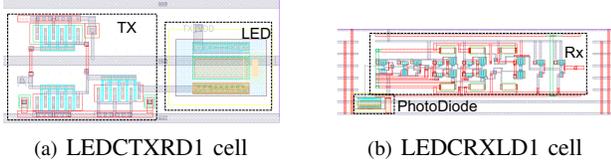


Fig. 3: Layout of standard cells in OESC library.

RF sense, they do experience possible cross-talk if co-parallel waveguides are too close to each other. Practically this sets a limit on the pitch of waveguides that can be used. In this work, the pitch was limited by the standard cell height of $5 \mu\text{m}$ (one waveguide per standard cell) for the $0.18\mu\text{m}$ CMOS node rather than cross-talk concerns. Second, unlike electrical interconnects, optical interconnects can theoretically cross each other, though accurate optical modeling would be needed to design the optical junction such that the signal loss and cross-coupling at that point does not cause problems in data transmission. Third, unlike electrical interconnects, optical signal propagation through tight bends (including out-of-plane propagation) is highly lossy relative to the almost optically-lossless propagation in a straight line. The upshot of this is that optical links are best formed as long, straight point-to-point links.

For the purposes of this work, we constrained our place and route tool to **not allow** waveguide crossings and bends on a plane, as a simplification to enhance the robustness of the findings from this initial study. We use two optical planes - one for horizontal waveguides, and one for vertical.

C. Electronic CMOS Circuits Design

Fig. 2 shows the schematic of our LED link. A CMOS transmitter (Tx) drives the InGaP LED, which is optically connected to the PD at the receiving end, through the waveguide. A CMOS receiver (Rx) senses the current from the PD and converts that to an output voltage. For each data rate, we target an end-to-end delay of 1-cycle, and the same bandwidth as an electrical link. We design a library of Tx and Rx circuits in $0.18\mu\text{m}$ CMOS targeting both conservative and aggressive optical device models across different data rates. The LED and the PD used a supply voltage of 4V each. The CMOS Tx and Rx use a 1.8V supply (set by the $0.18\mu\text{m}$ process).

Tx. The LEDs are sized to meet the bias requirements at the given frequency of operation. Device reliability is one key concern while driving an LED at large bias voltages as the CMOS devices and the metal routing should be designed to carry the DC current through the photonic devices, as well as ensuring that the CMOS junction voltages are not exceeded when LED is turned ON. We design our CMOS Tx in an open-drain architecture. The ON-resistance of the driver is designed

TABLE I: Area and Energy of O-E circuit components

Data (Mbps)	Area ($\mu\text{m} \times \mu\text{m}$)				Energy (fJ/bit)			
	LED	PD	Tx	Rx	LED	PD	Tx	Rx
1000	1×10	1×15	27×30	65×25	49	8	37	360
500	1×8	1×12	27×30	65×25	49	8	37	360
250	1×5	1×10	27×27	60×25	41	8	34	360
10	1×2	1×2	20×20	55×25	40	8	28	360

considering the LED on-resistance while keeping the power supply consumption from the LED supply under check. The LED capacitance was empirically calculated as a few fFs and this did not affect the loading of the driver pad.

Rx. The CMOS receiver (Rx) uses a two-pole, three-stage Trans Impedance Amplifier (TIA) to mitigate trade offs with signal, noise and bandwidth. The size of the TIA is kept minimal by replacing the standard n-well or poly resistors with the MOS devices operating in the triode region. The PD (reverse-biased LED) current is of the order of just a few $n\text{A}$ and the TIA is designed to generate a 50-75mV swing from it. One of the challenges of the Rx design is the need for a differential receiver owing to the small swing at the TIA output. A fully differential receiver will need an additional optical link running in parallel, increasing the silicon area. We chose a pseudo differential receiver design, whereby a fixed bias voltage, based on the TIA output voltage, is used as one input to the differential comparator. This bias voltage can be generated on-chip from stable voltage regulator outputs. This brings in a 50% area improvement and power benefits against traditional optical link architectures using fully differential receiver designs. Two gain stages are used to boost the differential amplifier output before feeding it to a differential-to-single ended (D2S) converter. The D2S output followed by a buffer yields the receiver output at the given data rate with minimal jitter.

Since our optical link is an intra chip design, ESD circuits are obviated at the transmitter and receiver side. In addition, the receiver circuits also obviates any secondary ESD protection, like Charge Device Model (CDM) circuitry. These two factors together with absence of bond pads, help reduce the parasitic capacitance seen by both the Tx and the Rx.

Table I lists the area and energy breakdowns of our designed opto-electronic components with *conservative* LED/PD models, as a function of data rates. These results are from Virtuoso simulations. We pick the 1Gbps designs in our evaluations. The CMOS Rx dominates the energy and area. The *aggressive* LED/PD device models in the PDK - which have not been validated by fabrication yet - have much lower turn on voltage, and an order of magnitude higher current. This helps reduce the supply voltage for the LED bias considerably, guaranteeing reliable CMOS junction voltages for the Tx. In addition, a higher current from the PD is available, which reduces the area and complexity of our CMOS Rx cells by up to 60%.

D. Standard Cell Library

We build an opto-electronic standard cell (OESC) library behind which the E-to-O and O-to-E complexity is hidden from designers. We provide cells for both the conservative and the aggressive device models.

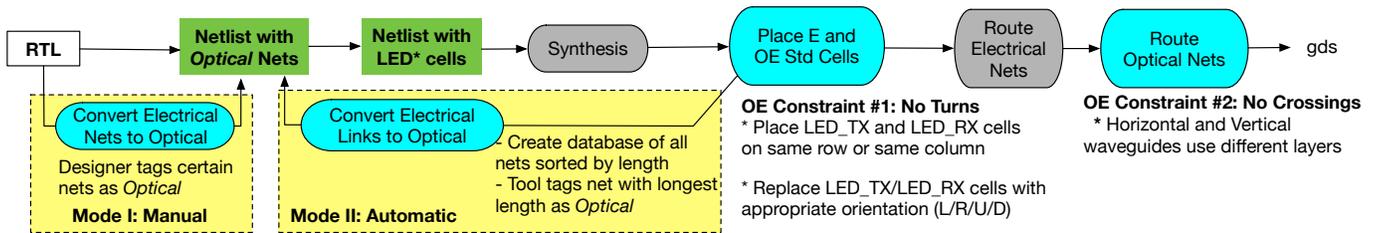


Fig. 4: CAD Flow. LEDs are added either manually or automatically via an iterative process. The place-and-route is fully-automated.

LED_TX is a group of OESCs performing E-to-O transmission (input voltage to output photons). Each LED_TX cell has an input port to send signals via a CMOS Tx to a LED. Photons are generated in the LED and dispersed from the cell through the optical port of the cell. A portion of particle light is channeled through the waveguide with minimum loss to a LED_RX cell. Fig. 3(a) shows one of the LED_TX cells, LEDCTXRD1, in which the electrical input and optical output ports are at the left and the right of the cell, respectively.

LED_RX is a group of OESCs for O-to-E transmission. Each LED_RX cell has an optical input port to detect photons from the optical output port of a LED_TX cell via a waveguide. The photons are transformed to an output current by the photo diode which is then converted to a voltage by the CMOS Rx. The voltage swing is about 300 mV, and thus the Rx uses an amplifier, with a static reference voltage input, to produce a robust eye diagram. Fig. 3(b) shows one of the LED_RX cells in which the optical input and electrical output ports are at the left and the right of the cell, respectively.

Cell Orientation. In electrical cells, nets connecting to the cell can enter from any of the four directions. In LED cells, however, we require four variants depending on the direction (Left/Right/Up/Down) of the waveguide to transmit (or receive) optical light. Since we use two optical planes, the Left and Right cells are used on the horizontal plane, while the Up and Down cells are used on the vertical plane.

We have a total of 16 cells in our OESC. We did not design additional cell instances based on waveguide lengths (analogous to small and large sized CMOS cells based on drive strength) since the minimum sized ones can easily drive a waveguide up to 10mm. For each standard cell, we provide a .lib/.db file for timing.

We add the **waveguide** information exactly like electronic metal layers into a foundry CMOS Library Exchange Format (lef) technology file. We also add the OESC cells like CMOS standard cells into a foundry standard cell “lef” file. This method allows the lef files to be used just like CMOS standard cells and electronic metal layers by the place-and-route (P&R) tool. Two waveguides are not allowed to cross each other, just like two nets on the same metal layer cannot cross each other. We have two waveguide layers (Section II-B), one for horizontal and one for vertical waveguides. Our methodology is general to handle any number of layers.

III. AUTOMATED CAD TOOL SUPPORT

Fig. 4 shows our CAD flow which can be integrated into any CAD tool; we used a combination of Synopsys and Cadence

tools with a commercial 0.18um CMOS PDK + 0.25um LEES PDK [6], [25].

LED Insertion. We add LED links by introducing the LED_TX and LED_RX cells into the netlist. This can be done via two modes: Manual or Automatic. The place-and-route flow after this takes care of automatically placing these and routing the waveguides subject to optical device constraints.

Mode I: Manual. We start with the RTL for the chip, and let the designer manually tag certain nets as *optical* by adding a special keyword in front of the desired nets. Examples could be long links between cores, or from a core to a cache. These nets are converted to LED links by the synthesis to layout flow described below.

Mode II: Automatic. Post placement, a database of all nets are created, sorted by their length. This is done pre-routing, since the routing step traditionally breaks long nets into a series of shorter metal wire segments connected by vias. If the net with the longest length is greater than the crossover point (which is an input parameter, e.g., 2.5 mm from Fig. 1(c)), it is tagged with the keyword *optical* in the netlist, and sent through placement again. *This is done iteratively till the placer can no longer place all LED cells within the given area subject to the optical constraints of no turns and no crossings.*

The two LED insertion modes can be enabled separately or together for design-space exploration.

Synthesis. A parser parses the RTL and adds new modules called LED_TX and LED_RX before and after the *optical* nets respectively, and then performs traditional timing synthesis, treating these cells as black boxes. As described earlier in Section II-C, the LED links are designed targeting the same delay as electrical links.

Placement. The placement follows the traditional constraints of area and timing. In addition, we introduce additional constraints for the OE cells, namely *no turns*. We guarantee this by developing an algorithm to place a pair of connected LED_TX and LED_RX cells in the same row or in the same column. The corresponding waveguides would be horizontal or vertical respectively. Additionally, no other LED cells should be placed between them. The algorithm makes multiple passes to produce a placement that can place all LED cells in the design within the target area. The placer also replaces the LED_TX/LED_RX cells with actual cells from the library with the correct orientation (Left/Right/Up/Down). If the placement fails, the nets that could not be made optical are converted back to electrical nets by converting the LED_TX and LED_RX cells to BUF or INV from the electrical standard cell library.

Routing. So far the P&R tool has not distinguished between

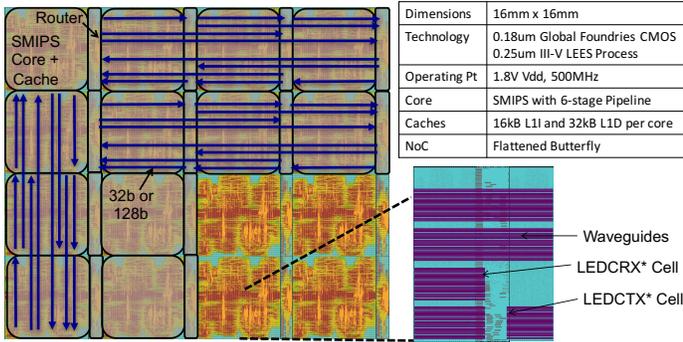


Fig. 5: Case Study: 4x4 Multicore with Flattened-Butterfly Network-on-Chip (NoC).

the electronic and the opto-electronic cells. Once placement is complete, we first parse the placement file to collect all the LED* cells, and disable routing between the LED_TX and LED_RX pairs. We now perform routing. This routes all the electrical nets, i.e., the tool routes between all the electronic cells, and between the electronic and opto-electronic cells using the PDK's metal layers and its own optimized algorithms. Next, we constrain the tool to only route between the LED_TX and LED_RX pairs using the waveguide layers. The placement step already guaranteed no turns. The routing steps guarantees the second constraint, *no crossings*, by using two waveguide layers to route the horizontal and vertical waveguides, respectively.

We leverage the .lib/.db files of our standard cells to perform timing closure analysis just like a conventional design.

IV. CASE STUDIES

We present two case studies with our CAD tool, demonstrating both the feasibility and viability of replacing electrical links both, long global/semi-global and short local. We target 1 Gbps (1cycle@1GHz) for both LED and electrical links.

A. Core-to-Core LED interconnects

We built a 16-core chip with SMIPS processors [4] connected via a modified version of the Flattened Butterfly [14] network-on-chip (NoC) as shown in Fig. 5. We chose this topology as it uses long links (for low hop counts) which are attractive candidates for replacement with optical links [8]. Each router has a dedicated link to every other router in its row and column. We run our tool in Mode I (designer-driven) and replace these with horizontal and vertical LED links respectively using two waveguide layers. We use electrical links for the local link between the router and the core (as those are less than 1mm), and LED links between routers, as those are all 4mm and longer, guided by Fig. 1(c).

In the conservative device models, the area of the LED_RX cells is large (Section II-C) and sets a limit on the number of waveguides that can fit within a tile dimension. We use a 32-bit datapath. i.e., 6x32 horizontal and vertical waveguides at each router as shown in Fig. 5, with the conservative-model based standard cells, and a 128-bit datapath with the aggressive ones.

We ran RTL simulations of different traffic patterns through the NoC and fed the router and link activities through the

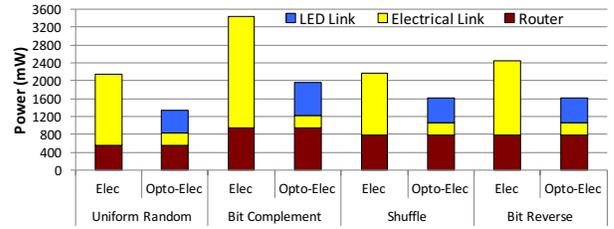


Fig. 6: Power Breakdown of 4x4 NoC with 128b datapath for different traffic patterns.

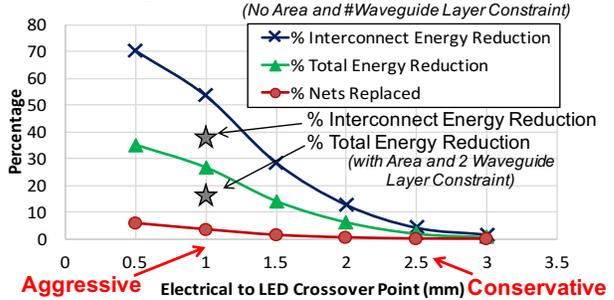


Fig. 7: Energy Reduction with LED links inside a core.

extracted netlist to compute the power consumption of the NoC in both designs. With the conservative area-limited cells (32b), we observed 10-21% reduction in network power across the patterns, while with the aggressive area-optimized cells (128b datapath), the reduction increased to 21-39% (Fig. 6). Bit Complement shows the most benefit as it has the heaviest cross-chip traffic.

The overhead of LED links comes in the form of area, since the OESC cells are an order of magnitude larger than electrical cells (due to E-to-O and O-to-E circuitry), and increase router area by 2X (conservative) and 1.25X (aggressive), as shown in Table II. However, since the core dominates the area of a tile, the overall area overhead is just 2.4-5.6%.

The key takeaway message is that our conservative and aggressive optical LEDs can enable designers to get up to 39% energy savings in the on-chip interconnect network, at the cost of up to 5% chip area overhead. And this comes at minimal design cost due to our automated tool.

This case study targeted the NoC energy but did not affect the energy of the core, which was found to be about 5X that of the NoC. The next case study addresses that.

TABLE II: Standard Cell Area Overheads (Per Tile)

	Core	Cons. (32b)		Aggr. (128b)	
		Rtr	LED	Rtr	LED
Num Cells	324073	22674	240	37767	960
Area (μm^2)	8360174	434544	523440	945087	236160
% of Total	~89	4.6	5.6	9.9	2.4

B. Within-core LED interconnects

Traditionally the role of on-chip photonics has been seen to be for replacing global/semi-global links between cores [8], as demonstrated above. However, we believe that LED links offer a unique opportunity to potentially replace links within a core as well. We ran our CAD tool in Mode II (Section III), letting it replace as many links as possible with optical links.

First, we limit the design to use two waveguide layers, and fit within the same area as the electrical design, and use the

aggressive (low area) device models. This leads to an energy savings of 17% in the core as shown in Fig. 7.

Next, we run a limit study, where we (i) let the tool use any number of waveguide layers as it needed, (ii) remove area constraints, and (iii) sweep the crossover point at which LED links surpass electrical links in terms of energy consumption. In this study, all nets that are longer than the crossover length are replaced with LED links by our tool. Fig. 7 plots our results. We found only 5% of the nets to be longer than 1mm inside the core, but replacing these with the aggressive model LED links results in 53% (i.e., more than 2X) reduction in interconnect energy due to the fixed energy cost irrespective of routing distance, which translates to 27% reduction in the total core energy. With a crossover point of 0.5mm, the savings can go up to 70% and 35% respectively.

This limit study shows the potential of LED links in future SoCs. SoCs are expected to be more power-limited than area-limited; an area cost of the OE standard-cells could be an acceptable trade-off for the energy-savings they provide. Adding more waveguide layers, in lieu of electrical metal layers may also be an effective trade-off. *Active research in low-power and low-area LED devices can provide a solution for lowering the energy consumption of cores for the extremely energy-constrained chips for the mobile and IoT domain.*

V. RELATED WORK

Commercial CAD tools have recently started offering support for photonics integrated circuits, such as Synopsys RSoft Optsim [2] and Cadence EPDA [3]. However, the support is for custom circuit design rather than automated photonic link generation since the focus of these tools is on off-die IOs, where today's electrical high-speed links are also custom designed to match specific requirements.

Academic research in CAD for photonics has been burgeoning to optimize the placement and routing of photonics devices, such as the minimizing and tuning of laser power and waveguide crossings and bends [5], [9], [12], [18], [21], [24]. In contrast, we target novel on-die LEDs as the light source, focusing on a standard-cell library of opto-electronic cells and their automated placement. In the future, when the monolithic-integrated CMOS+III-V LEES process scales to enable more photonics layers to be integrated, more sophisticated place-and-route algorithms such as those proposed by these prior works will be able to further improve the energy and area footprint of such LED-driven optical wires.

Design automation for on-die photonics has been looked into from synthesis to GDS in academia, but the focus lies in using integrated optics for logic [7], not just as optical interconnects for on-die communications. Hence, these prior works start from logic synthesis. By leveraging the LEES process we can continue to rely on CMOS for logic, and just leverage optics for what it does best: communications.

To the best of our knowledge, this paper is the first to demonstrate an opto-electronic full link that can replace copper wires, and present a synthesis-to-fabrication tool flow that works within a state-of-the-art commercial toolchain.

VI. CONCLUSION

In this work, we make a case for using novel μ -LED-based optical interconnects on-chip as direct replacements for electrical links, both across cores and within cores, leveraging a monolithically integrated CMOS + III-V process called LEES and measured data from InGaP LEDs. Working across the materials, devices, circuits, CAD, and architecture stacks, we present a tool flow that hides the complexity of optical devices and associated electrical circuitry behind standard cells, and automatically replaces electrical nets by LED links subject to area constraints. We believe this work opens up a plethora of cross-layer research opportunities.

REFERENCES

- [1] GPU Computing to Exascale and Beyond, Bill Dally, Keynote SC'2010.
- [2] <http://news.synopsys.com/index.php?item=123417>.
- [3] www.cadence.com/solutions/photonics/Pages/default.aspx.
- [4] Arvind, R. S. Nikhil, J. Emer, and M. Vijayaraghavan. *Computer Architecture: A Constructive Approach*. MIT, 2012.
- [5] A. Boos *et al.* Proton: An automatic place-and-route tool for optical networks-on-chip. In *ICCAD*, 2013.
- [6] S. B. Chiah *et al.* A Hybrid Process Design Kit: Towards Integrating CMOS and III-V Devices. In *Workshop on Compact Modeling*, 2016.
- [7] C. Condrat, P. Kalla, and S. Blair. Thermal-aware synthesis of integrated photonic ring resonators. In *ICCAD*, 2014.
- [8] Y. Demir and N. Hardavellas. SLAC: Stage laser control for a flattened butterfly network. In *HPCA*, 2016.
- [9] D. Ding *et al.* O-router: an optical routing framework for low power on-chip silicon nano-photonics integration. In *DAC*. ACM, 2009.
- [10] E. Fitzgerald *et al.* Enabling the integrated circuits of the future. In *IEEE International Conference on Electron Devices and Solid-State Circuits (EDSSC)*, pages 1–4, 2015.
- [11] A. Gorin *et al.* Fabrication of silicon nitride waveguides for visible-light using PECVD: a study of the effect of plasma frequency on optical properties. *Optics Express*, 16(18):13509–13516, 2008.
- [12] G. Hendry, J. Chan, L. P. Carloni, and K. Bergman. VANDAL: A tool for the design specification of nanophotonic networks. In *DATE*, 2011.
- [13] A. Kelly *et al.* High-speed gan micro-led arrays for data communications. In *IEEE International Conference on Transparent Optical Networks (ICTON)*, 2012.
- [14] J. Kim *et al.* Flattened butterfly topology for on-chip networks. In *MICRO*, pages 172–182, 2007.
- [15] K. Lee *et al.* Monolithic integration of III–V HEMT and Si-CMOS through TSV-less 3D wafer stacking. In *IEEE Electronic Components and Technology Conference (ECTC)*, 2015.
- [16] O. López *et al.* Highly-efficient fully resonant vertical couplers for inp active-passive monolithic integration using vertically phase matched waveguides. *Optics express*, 21(19):22717–22727, 2013.
- [17] D. Miller. Device Requirements for Optical Interconnects to CMOS Silicon Chips. *Photonics in Switching*, 2010.
- [18] J. R. Minz, S. Thyagara, and S. K. Lim. Optical routing for 3-d system-on-package. *IEEE Transactions on Components and Packaging Technologies*, 30(4):805–812, 2007.
- [19] J. Piprek *et al.* Physics of waveguide photodetectors with integrated amplification. In *Integrated Optoelectronics Devices*, pages 214–221. International Society for Optics and Photonics, 2003.
- [20] G. Roelkens *et al.* III-V/silicon photonics for on-chip and intra-chip optical interconnects. *Laser & Photonics Reviews*, 4(6):751–779, 2010.
- [21] C.-S. Seo *et al.* Physical design of optoelectronic system-on-a-package: a cad tool and algorithms. In *ISQED*, 2005.
- [22] A. Shacham, K. Bergman, and L. P. Carloni. Photonic networks-on-chip for future generations of chip multiprocessors. *IEEE Transactions on Computers*, 57(9):1246–1260, 2008.
- [23] C. Sun *et al.* Single-chip microprocessor that communicates directly using light. *Nature*, 528(7583):534–538, 2015.
- [24] A. von Beuningen and U. Schlichtmann. PLATON: A Force-Directed Placement Algorithm for 3D Optical Networks-on-Chip. In *International Symposium on Physical Design*. ACM, 2016.
- [25] B. Wang *et al.* On-chip Optical Interconnects using InGaN Light-Emitting Diodes Integrated with Si-CMOS. In *Asia Communications and Photonics Conference*. Optical Society of America, 2014.