Optical Routing for 3-D System-On-Package

Jacob Rajkumar Minz, *Student Member, IEEE*, Somaskanda Thyagara, *Student Member, IEEE*, and Sung Kyu Lim, *Senior Member, IEEE*

***Abstract*—In this paper, we present the first optical router for 3-D system-on-package (SOP). Recent advances in optical device integration for SOP offer drastic advantages over electrical interconnects. We propose efficient algorithms for the construction of timing and congestion-driven waveguides taking into account the optical resource constraints. Our experimental results suggest that smart placement of waveguides coupled with other routing techniques can reduce electrical wirelength by 11% and improve performance by 23%, when a single optical layer is introduced for every placement layer.**

***Index Terms*—Optical routing, system-on-package (SOP), waveguide construction.**

I. Introduction

THE system-on-chip (SOC) paradigm is a new system integration approach, where not only more and more transistors but various mixed-signal active and passive components are also integrated into a single chip. However, the systems community is beginning to realize that SOC presents fundamental, engineering, and investment limits [1]. This led to the 3-D system-on-package (SOP), where the package not the chip becomes the medium for system integration. An illustration is shown in Fig. 1. This improvement comes in two ways: 1) It uses CMOS-based silicon for what it is good for, namely, for transistor integration, and the package, for what it is good for, namely, radio frequency (RF), optical, and digital integration by means of integrated circuit (IC)-package-system co-design. The SOP package, therefore, overcomes both the computing limitations and integration limitations of SOC, system-in-package (SIP), multichip module (MCM), and traditional system packaging.

Optoelectronics, which today finds use primarily in the back planes and in high-speed board interconnects, is fast moving onto SOP as chip-to-chip for high input/output (I/O) and high-speed interconnections. Optical interconnects are replacing copper ones and thus addressing both the resistance and crosstalk issues of electronic ICs [2]. Wide area, high speed optical clock, and data transport simplifies the digital architecture because fewer parallel transmission lines are needed for the same bandwidth. Also, optical links have low crosstalk and are not susceptible to electromagnetic interference (EMI) noise,

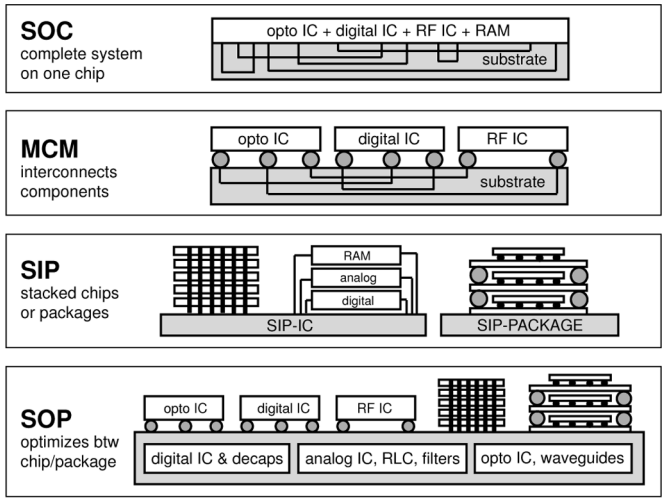

Fig. 1. Comparison among SOC, MCM, SIP, and SOP.

thus also reducing the need for decoupling capacitors. In addition, wavelength division multiplexing (WDM) allows a single waveguide to be shared among multiple optical interconnects, thereby significantly reducing the area used by an interconnect.

Chang *et al.* [2] recently proposed a low-cost opto/digital integrated circuit on a standard printed wiring board. They successfully developed a low-temperature polymer process for fabricating and integrating optoelectronic components such as lasers, waveguides, and photo-detectors onto printed wiring boards for mixed signal SOP applications. Chen *et al.* demonstrated a fully embedded implementation of high-speed optical communications within one board [3]. A 12-channel linear array of thin-film polyimide waveguides, vertical-cavity surface-emitting lasers and silicon photodetectors were used for this experiment. Also, a 1-to-48 optical clock signal distribution network for Cray T-90 supercomputer was demonstrated. The WDM concept has been experimentally demonstrated in [4]. A pseudo-random data generator circuit on the chip was used to feed 10 channels with different data that were combined at a reflective grating and transmitted through the optical fiber. GaAs diodes that were flip-chip bonded onto 0.5 μm silicon CMOS were used as modulators as well as detectors to spectrally split the incoming beam by wavelength. Such experiments have proved that optical routing on the package level is a promising solution.

In this paper, we present the first optical router for 3-D SOP. A recent work on placement for 2-D SOP is presented in [5]. The authors in [6] performs clock routing using optical waveguides. Our approach is to start with a pure electrical interconnect routing solution, build a set of customized waveguides, and convert a subset of electrical wires into optical interconnects by

Manuscript received December 21, 2005; revised September 3, 2006. This work was supported by MARCO C2S2 Grants and NSF CAREER under Grant CCF-0546382. This work was recommended for publication by Associate Editor F. Shi upon evaluation of the reviewers comments.

J. R. Minz is with the Synopsys Corporation, Sunnyvale, CA 94089 USA.

S. Thyagara is with Advanced Micro Devices, Inc., Austin, TX 78758 USA.

S. K. Lim is with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332 USA.

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TCAPT.2007.906300

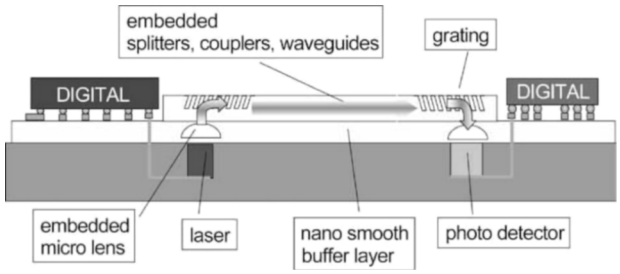

Fig. 2. Optical routing in SOP [2].

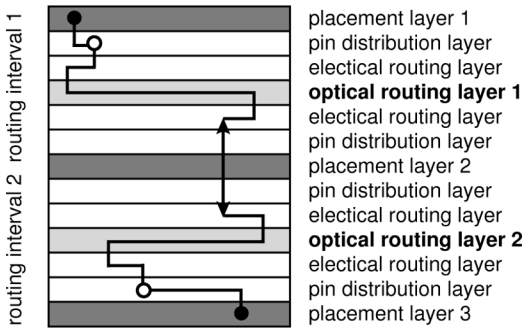

Fig. 3. Illustration of the layer structure and routing resource in SOP. The black and white dots, respectively, denote the original and redistributed pins. The arrowed line denotes feed-through via.

rerouting them to use these waveguides. The WDM capability allows us to assign multiple electrical nets to each waveguide, where each waveguide has a limit on the number of different wavelengths it can handle. In addition, various kinds of modules such as lasers, splitters, couplers, and photo detectors used for each waveguide occupy physical space in the layout. Thus, our optimization goal is to minimize the total wirelength and delays under various layout capacity constraints.

This paper is organized as follows: Section II presents an overview of the optical routing technology for SOP. Section III presents the problem formulation. Section IV presents our 3-D SOP optical routing algorithm. Experimental results are shown in Section V, and we conclude in Section VI.

II. SOP Optical Interconnect Technology

In Fig. 2, are shown the basic technologies required for a fully integrated digital-optical micro-system. The text boxes indicate enabling integration technologies which have been developed to achieve full digital-optical functionality. Each high-frequency output port from a processor modulates a specific laser in an array. The digitized optical signal is coupled through a micro-lens array, into the optical signal distribution network comprising waveguides, splitters, couplers, gratings, etc., where it is transported to its destination. The optical signal is detected by a specific photodiode in an array of optical receivers, and converted to an electrical signal that is input into a specific port of the receiving processor. The signaling is bidirectional and nonblocking. Optical signals are coupled in and out of the optical transport network by a number of means that include gratings, lenses, waveguide end-mirrors, directional couplers and evanescent coupling. The entire opto/digital microsystems is built directly on the buffer layer which is fabricated on low cost FR-4 and APPE boards.

There are different types of signal losses in optical routing. First, certain optical devices have internal signal losses associated with them. For example, the splitting loss of an optical splitter is about 0.4 dB per splitter as demonstrated in [7]. Optical waveguides have internal as well as external losses. Internal losses are primarily due to material absorption and propagation. Absorption effects in general are more prominent at the higher wavelengths while propagation losses in waveguides are estimated at 0.36 dB/cm at a wavelength of 1.3 μm [8]. The length of a waveguide is critical from a physical design standpoint as a greater length corresponds to greater propagation loss. Extrinsic losses typically result from scattering off of profile roughness and defects. Second, signal losses also result from optical waveguide bending. For example, when a signal passes through a bent optical waveguide, the tangential velocity of the signal in a cladding layer will exceed the velocity of light. Hence this portion cannot stay in phase and splits away from the guided mode signal resulting in signal loss. As a result, a greater bending radius corresponds to a lower signal loss. Thus, we impose a limit on the maximum waveguide length and the number of bends while constructing waveguides in order to minimize signal loss.

III. Preliminaries

A. Problem Formulation

The layer structure in multilayer SOP is illustrated in Fig. 3. The placement layers contain the blocks (such as ICs, embedded passives, opto-electric components, etc), which from the point of view of physical design are just rectangular blocks with pins along the boundary. The interval between two adjacent placement layers is called the routing interval. A routing interval contains a set of routing layers sandwiched between pin distribution layers. At the core of the routing layer set is a single optical routing layer that contains both waveguides and optical modules such as laser, lens, detector, splitter, and grating. The pin distribution layers in each routing interval are used to evenly distribute pins from the nets that are assigned to this interval. Then these evenly distributed pins are connected using the routing layer pairs. A feed-through via is used to connect two pin distribution layers from different routing intervals. The nets which have all their terminals in the same placement layer are called i-nets, while the ones having terminals in different placement layers are x-nets.

Given a 3-D SOP placement, a set of nets N and number of placement layers L_p, the 3-D SOP Optical Routing Problem is defined formally as follows: generate a routing topology for each net $n \in N$, assign n to optical or electrical routing layer or both, and assign all pins of n to legal locations. All conflicting nets are assigned to a different routing layer. The goal is to minimize a weighted sum of: 1) the maximum delay among all nets; 2) total wirelength of all nets; and 3) total routing layers (electrical and optical) used. The constraints include via and wire capacity for each routing region in a given 3-D SOP structure.

B. Delay Model

Our approach is to first start with a pure electrical routing solution and build optical waveguides based on the usage of the

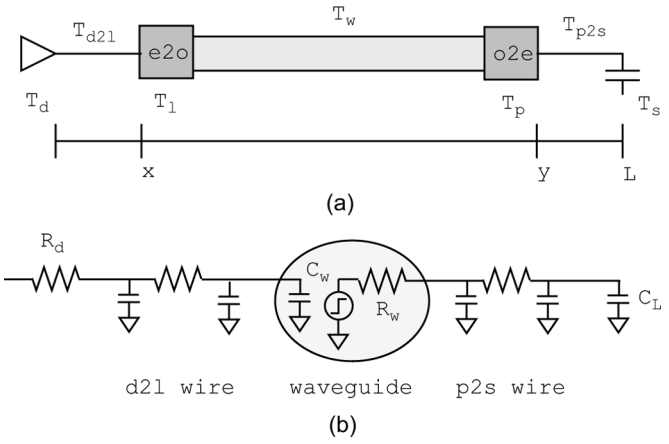

Fig. 4. (a) interconnect with both electrical and optical components and (b) RC-model, where C_w models the electrical-to-optical signal convertor and R_w models the optical-to-electrical signal convertor.

routing channels. We then convert a subset of these electrical nets to optical by utilizing the pre-built waveguides. Note that this electrical-to-optical conversion requires lasers and detectors. We model the following components in our electro-optical delay computation (see Fig. 4):

- T_d: delay of the driving gate;
- T_{d2l}: electrical wire delay from the driver to laser;
- T_l: intrinsic delay of the laser (= electrical to optical signal conversion);
- T_w: flight time of signal in an optical waveguide;
- T_p: intrinsic delay of the detector (= optical to electrical signal conversion);
- T_{p2s}: electrical wire delay from the detector to sink gate;
- T_s: delay of the sink gate.

The computation of electrical part ($= T_d, T_{d2l}, T_{p2s}$, and T_s) is based on Elmore delay model [9]

$$D_{\text{elec}}(s,v) = \sum_{k \in P} R_k \cdot C_k \tag{1}$$

where P is the path from source node s to sink v, R_k is the local resistance seen at node k, C_k is the downstream capacitance seen at node k. The computation of optical part ($= T_l, T_w$, and T_p) is based on linear delay model [6]

$$T_w = b \times l \tag{2}$$

where l is the length of the waveguide, and b is the constant representing delay per unit length in the optical interconnect. We use the values reported in [2] for T_l and T_p. Thus, the delay of any source-to-sink interconnect is the summation of all seven components mentioned above.

We use the following closed-form formula for the entire electro-optical interconnect (see Fig. 4):

$$\begin{aligned} D(x,y) = {} & R_d(cx + C_w) + rx\left(\frac{cx}{2} + C_w\right) \\ & + R_w(c(L-y) + C_L) \\ & + r(L-y)\left(\frac{c}{2}(L-y) + C_L\right) + t_d \\ & + \frac{1}{v_w}(y-x) \end{aligned} \tag{3}$$

where x, y are the start and end points of the waveguide, R_d is the driver resistance, C_L is the load capacitance, r, c is the unit wire resistance and capacitance, R_w is the opto-device output resistance, C_w is the opto-device input capacitance, t_d is the opto-devices intrinsic delay, v_w is the waveguide signal velocity, and L is the length of the wire. The convertors are modeled as buffers, and hence are assumed to have an input capacitance (for electrical-to-opto convertor) and output driver resistance (for opto-to-electrical convertor). The convertors also have intrinsic delay terms captured by t_d.

IV. SOP Optical Routing Algorithm

We develop the following algorithms for various steps required during the SOP optical routing process: 1) construction of optical waveguides based on performance and congestion consideration; 2) optimum net-to-waveguide mapping using maximum-flow minimum-cost network flow model; and 3) ripping and rerouting of existing nets to relieve congestion and make maximum utilization of waveguides.

A. Overview of the Algorithm

Compared with the conventional "single device multiple routing layer" model as in IC, PCB, and MCM routing, SOP routing is more general in that it requires "multiple device multiple routing layer" model. Due to the complexity involved in this 3-D routing process, we divide the entire process into the following steps.

1) Pin redistribution: we first determine which set of net segments are assigned to each routing interval. The pins from these nets are then evenly distributed in the pin distribution layers.
2) Topology generation: Steiner trees are generated for all nets in each routing interval so that the performance of the routed design is optimized.
3) Waveguide construction: a planar set of optical waveguides is constructed based on the topology generation in each routing interval. Each waveguide has up to a single bend and its maximum length is constrained.
4) Optical net selection: a subset of nets that makes the best use of the waveguides is chosen. Each net is assigned to a unique waveguide based on its proximity and capacity.
5) Optical routing: each optical net is rerouted so that it uses the entire waveguide it is assigned to while minimizing the length of its electrical portion.
6) Layer assignment: the routed nets are assigned to a unique routing pair in the routing layer so that the total number of layers used is minimized.
7) Local routing: the location of feed-through via for each net in the routing channel is determined. In addition, we finish the connection between the original and distributed pins.

We use the pin redistribution and layer assignment algorithms in [10] and topology generation algorithm in [11]. In addition, we use the congestion-driven rip-up-and-reroute algorithm in [12] for local routing. Therefore, the focus of this paper is to develop heuristics for the waveguide construction, optical net selection, and optical routing.

Timing-driven Waveguide Construction
input: electrical net routing results
output: timing-driven waveguide topology

```
1:  W = ∅;
2:  N = selected critical nets;
3:  sort N by criticality;
4:  while (N is not empty)
5:      B = set of segmented root-to-sink paths;
6:      for each (b ∈ B)
7:          L = candidate waveguide locations;
8:          l_i = minimum-cost waveguide location;
9:          P = constrained length waveguide at l_i;
10:         add P to the set of waveguides W;
11:     update waveguide usages;
12: return W;
```

Fig. 5. Pseudocode for timing-driven waveguide construction.

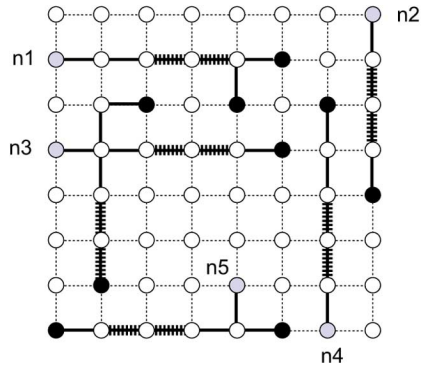

Fig. 6. Timing-driven waveguides for five timing critical nets. The net n_1, n_3, and n_5 are multi-pin nets. The gray and black nodes, respectively, denote the source and sink nodes. The waveguides are shown in dotted lines.

B. Waveguide Construction Algorithm

Our goal during optical waveguide construction is to choose the most optimum nets for optical routing based on their locations and properties such as wirelength, delays, and number of bends. In our model, each routing layer is represented by a grid graph $G(V, E)$, where each node specifies a region in the layer and each edge represents the adjacency between neighboring regions. The nodes and edges are characterized by capacities. Specifically, the nodes are annotated with optical device capacities and the edge capacities correspond to number of waveguides. A simple calculation based on the dimensions given in [2] for a 1 cm by 1 cm package, represented by a 10×10 routing grid, gives the node capacity to be 4 and edge capacity can vary between 4 to 8 [4]. In this paper, we assume the edge capacity to be 4. In addition, each waveguide can carry up to 10 multiplexed signals. This means there are up to 40 nets that can be assigned to each routing edge in G.

1) Timing-Driven Waveguides: The construction of the timing-driven waveguides considers the timing critical nets. The aim of the waveguide construction is to primarily reduce the worst delays of the design and secondarily improve electrical wirelength and routing layers cost. The timing criticality of a net is defined as the ratio between the net delay and the worst delay among all nets in the routing layer. The delay values are available since nets are already routed using electrical routing resources. A net is said to be a critical net if its criticality exceeds a predefined criticality level.

The pseudocode of the timing-driven waveguide construction algorithm is shown in Fig. 5. Our timing-driven approach is based on the following important observation.

Lemma 1: Given a multipin net n routed with pure electrical wires, converting a part of n into optical will not degrade the performance if 1) the overall topology is maintained, and 2) the delay reduction caused by optical waveguides exceeds the delay penalty caused by optical modules.

Thus, our timing-driven approach is to preserve the routing topology of the timing critical nets and convert parts of the nets into optical only when the overall delay reduction is positive.

The algorithm starts with the set of pre-selected timing critical nets, which are sorted by their criticalities. The paths for each of the critical sinks (sinks with delays above a threshold) of the net are computed. These paths are segmented at its branch points, and the set B contains a set of no-branch segments.[1] The candidate locations for waveguides are determined for each of the segmented paths. The optimal delay location for the waveguide is achieved by differentiating $D(x, y)$ with respect to x and solving it for zero

$$\frac{\partial D(x,y)}{\partial x} = \left(cR_d - cR_w\frac{\partial y}{\partial x}\right) + rc\left(x - (L-y)\frac{\partial y}{\partial x}\right) + \left(rC_w - rC_L\frac{\partial y}{\partial x}\right) + \frac{1}{v_w}\left(\frac{\partial y}{\partial x} - 1\right). \quad (4)$$

The optimal delay location is close to the center of the line if $R_d \approx R_w$ and $C_L \approx C_w$ and the size of the waveguide is constrained (i.e., $y \approx x+$ constant). The cost of waveguide insertion at each of the locations (l_i) is calculated based on the number of bends and distance from the midpoint (mid_i) of the path. The node and edge usages are also taken into consideration for device and waveguide allocation. The cost of the location is given by

$$\mathrm{cost}(l_i) = \alpha \cdot |l_i - |_i| + \beta \cdot \mathrm{bends} + \gamma \cdot \frac{\mathrm{usage}}{\mathrm{capacity}} \quad (5)$$

where α, β, γ are the weights for the cost factors. In our experiments, α and γ was chosen to be 1, whereas β was chosen to be 2. The waveguide is constructed at the location with the minimum cost (see Fig. 6). It is possible that the waveguides may cross each other at some nodes. However, such waveguides can be split at those junction points, and opto-devices plus a local electrical connection can be used to bridge the broken optical connections. A small penalty is incurred at those points for opto-conversions and local electrical wiring. The waveguides constructed are therefore guaranteed to be planar.

2) Congestion-Driven Waveguides: The objectives in the design of congestion-driven waveguides are to reduce the costs of electrical wiring and routing layers while not compromising design performance. Heavily used edges are obtained by noting the number of nets that pass through each edge and choosing edges which surpass a certain threshold. In our customized waveguide construction, heavily used edges are noted in each routing interval and planar waveguides are constructed in these areas to

[1]We assume that our timing-driven optical waveguides allow up to a single bend but prohibits branching.

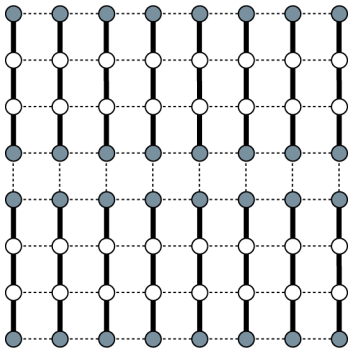

Fig. 7. Standard waveguides. The gray nodes indicate extry/exit points for each waveguide.

Congestion-driven Waveguide Construction
input: electrical net routing results
output: congestion-driven waveguide topology

```
1:  W = ∅;
2:  calculate routing edge usage;
3:  L = set of all heavily used edges;
4:  while (L is not empty)
5:      P = ∅;
6:      while (|P| < max_size)
7:          e = edge ∈ L located closest to P;
8:          if (adding e to P does not cause a bend)
9:              add e to P;
10:         fill gap between e and P with edges in L;
11:     add P to W if W is still planar;
12: return W;
```

Fig. 8. Pseudocode for congestion-driven waveguide construction.

reduce the number of electrical nets using these edges and thus reduce congestion.

Fig. 8 shows our customized waveguide construction algorithm. Our greedy approach starts with a set of heavily used routing edges L from a given global routing solution. Our "cluster growth" based method attempts to grow the current waveguide P by adding an edge e from L based on the distance between e and P. The waveguide has to be a straight line with no bends all the time[2], and its maximum length is also limited. In case there exists a discontinuity between e and P, we fill the gap using any intermediate edges to keep P a straight line. Upon the completion of constructing a single waveguide, we check for the planarity, i.e., we check to see if adding P to the current set of waveguide W does not cause any crossing. Last, the set W contains all legal and optimized waveguides. A minimum length of two edges and a maximum length of four edges is allowed for any waveguide in our experiment. This is because a single-edge waveguide is impractical from a cost standpoint, whereas waveguides with length greater than four edges would suffer from signal loss and low utilization.

3) Standard Waveguides: The method described for the congestion-driven waveguides requires a different layout for each optical layer that is generated for a specific routing interval.

[2]Note that we do not allow any bends in our congestion-driven waveguides unlike timing-driven waveguides. This is because the strict requirement reduces the signal loss significantly but has very little impact on the utilization under the congestion objective.

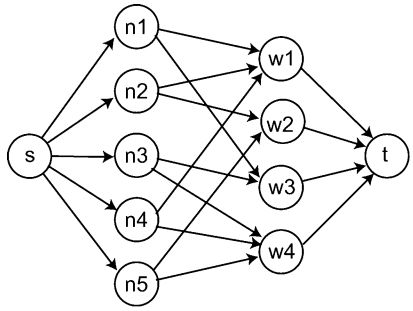

Fig. 9. Flow network for opto-net selection.

From a fabrication point of view, this would be more time consuming as there exists no standard design for any optical layer. In our *Standard Waveguide Construction*, we address this issue by choosing a standard subset of edges in every routing interval irrespective of their individual usages. This approach, however, can result in lower utilization of waveguides which makes it an inefficient solution. Our standard waveguide construction is straightforward in that a fixed subset of edges is used for the waveguides for every optical layer in the SOP. As shown in Fig. 7, each waveguide has a length of 4 edges and there are 16 waveguides in each routing layer. The waveguides could have also been constructed using horizontal edges, but that makes a negligible difference in the overall results.

C. Optical Net Selection Algorithm

The next step is to select a set of nets and assign them to the waveguides we constructed. We use a minimum-cost maximum-flow network flow model to select the most optimal nets to be assigned to each waveguide. The goal of this model is to minimize the overall cost while respecting the individual edge capacities. This model is illustrated in Fig. 9, where a source node connects to the net nodes (each net is assigned a node) and all waveguide nodes (each waveguide is assigned a node) are connected to a sink node. A net node is connected to a particular waveguide node if all the routing edges in that waveguide are contained in the net. Each edge in this flow graph is assigned a capacity and cost value. We assign a capacity of 40 and cost of 0 for every waveguide-to-sink edge because a maximum of 40 nets can use the waveguide and no waveguide costs different from another. The minimum length of nets considered for optical routing is 4 because it needs to utilize all edges of its waveguide. Each net-to-waveguide edge has a cost equal to number of edges in the net that are not in any waveguide because such edges will have to be routed electrically and it is better to minimize these connections in an optically routed net. Finally, the flow analysis provides flow values for each edge that tells us how useful it is to retain that edge in our final model. For example, when a waveguide has more than 40 nets mapped to it, we utilize these flow values to only retain nets with the 40 maximum flow values.

D. Optical Routing

The customized waveguides either cover timing critical net segments or heavily congested areas of the routing layer. Hence,

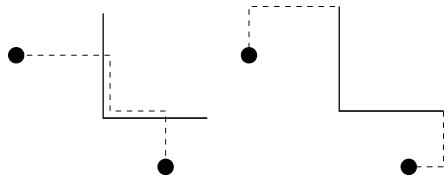

Fig. 10. Rerouting of the net to go through the optical channel. The figure shows electrical routes (dashed line) before and after rerouting through the optical waveguide (bold line).

after routing the net through the optical waveguides, it is desirable for the electrical part of the net not to go through these optical regions (see Fig. 10). In order to achieve this, all nets that utilize optical waveguides are segmented into three parts and rerouted: 1) the electrical part of the net topology containing the entry of the waveguide, 2) the optical waveguide, and 3) the electrical part of the net topology containing the exit point of the waveguide. Then our goal is to reroute 1) and 3) while not using 2). The nodes in the routing graph corresponding to the waveguides are deleted, ensuring that the net subtrees do not use them, thereby reducing overall electrical routing congestion. However, if the worst delay for the net deteriorates beyond the performance budget after optical rerouting, the net is converted back to electrical.

Rectilinear Steiner arborescence creates a routing tree with minimum source to sink distances for all sinks, and the weight of the tree is also minimum. It was shown that constructing minimal rectilinear Steiner arborescence is NP-complete [13]. We modify the RSA/G heuristic [11] to consider congestion-driven weights during the arborescence construction for simultaneous performance and congestion consideration. Our algorithm starts by constructing the weighted shortest path subgraph for a given net n to be rerouted by finding out all shortest paths from the source to sink in n. The weight is based on the current usage of the routing resource graph. The nodes in the subgraph are ranked according to the weighted distance from the source node (higher distance translates to higher ranks). The heuristic constructs the topology by considering the nodes in descending order of their ranks and merging them iteratively to form Steiner nodes. During the tree construction we do not allow merging at the routing nodes that are occupied by the existing waveguides.

Let s_n be the source of net n. T_n is the set of sink nodes of n. Let i_n and o_n be the entry and exit point of the waveguide that n contains. Then, I_n includes i_n as well as the set of all sink nodes of n located close to i_n. O_n is the set of all sink nodes of n located close to o_n. Thus, $I_n \subset T_n$ and $O_n \subset T_n$. Let X be the set of all nodes included in the waveguide of n except i_n and o_n. Finally, we construct the Steiner arborescence from s to I_n and from o_n to O_n while excluding X. This formulation can be easily extended to consider the nets with multiple waveguides.

V. Experimental Results

We implemented our algorithms in C++/STL and ran our experiments on Linux Beowulf clusters. For our experiments we used the standard GSRC floorplanning benchmarks. The blocks are placed into four-layer SOPs using our SOP floorplanner [10].[3] We used the technology parameters for 0.13 μ process [14] for Elmore delay computation. Specifically, the driver resistance (R_d) of 29.4 kΩ, input capacitance (C_L) of 0.050 fF, unit-length resistance (r) of 0.82 Ω/μm and unit-length capacitance (c) of 0.24 fF/μm are used. For calculating the optical delay, v_w was chosen to be 1.936 $\times 10^8$ m/s and t_d was assigned 200 ps [2]. The values of R_w and C_w was chosen to be the same as R_d and C_L. The wirelengths reported have been scaled down by 10^5 μm. All routing results are evaluated against the baseline pure-electrical routing. The runtimes reported are in seconds.

TABLE I
Standard versus Customized Waveguides

	Standard Waveguides				Customized Waveguides			
ckt	nets	WG	WL	Util	nets	WG	WL	Util
n30	55	60	212	8.83	224	41	740	54.63
n50	87	60	348	14.5	277	35	906	79.14
n100	123	60	492	20.5	259	27	889	95.92
n200	235	60	940	39.16	190	19	640	100.00
n300	218	60	872	36.33	358	36	1242	99.44

Table I shows the preliminary results obtained after the construction of waveguides using the two techniques (congestion-driven customized and standard). The two sets of waveguides were passed through the same opto-net selection algorithm and hence these results are solely based on the difference in waveguide construction. Wirelength savings is calculated as the amount of electrical wiring (in terms of grid edges) that has been converted into optical routing. Waveguide utilization is a ratio of number of nets actually using the waveguides to the maximum number of nets that could use these waveguides. It is evident that customized waveguides provide much better waveguide utilization and wirelength savings. Also, they use less number of optical waveguides but convert more electrical nets into optically routed nets.

Table II compares pure electrical routing with congestion-driven optical routing, optical routing using standard waveguides, and timing-driven optical routing. In all cases, only one optical layer is introduced per routing interval. An average wirelength saving of 8.5% was achieved using customized waveguides. The maximum wirelength improvement was 11% for n50. The amount of wirelength saved is less with the standard waveguides (3.7%). The number of layers reduce by one for most cases in congestion-driven optical routing but mostly remain the same for standard waveguides. The performances of the designs are preserved for all benchmarks, with only nominal improvement in some instances. The wirelength saving in timing-driven optical routing is 3.6% and the performance enhancement is 19% on the average.

Table III compares timing-driven optical routing with different criticality levels. Reducing criticality levels increases the number of critical nets considered for waveguide constructions. While the chances that delay improves increase, the probability of optical resources capacity violation increases. The wirelength savings using criticality levels of 0.90, 0.85, 0.80, are 1.5%, 2.6%, and 4.3%. The performance improvements

[3]Our attempt to compare to the routing results reported in [5] was not successful for several reasons. First, [5] reports only 2-D SOP results. Second, [5] only reports delay improvement in percentage and does not provide details on delay models used.

TABLE II
COMPARISON BETWEEN ELECTRICAL, CONGESTION-DRIVEN, STANDARD WAVEGUIDES AND TIMING-DRIVEN OPTICAL ROUTING WITH NO VIOLATIONS. WE REPORT THE TOTAL WIRELENGTH, ROUTING LAYER (XY) USAGE, AND MAXIMUM NET DELAY (NS)

	Electrical			Congestion Optical			Standard Optical			Timing Optical		
ckt	WL	lyrs	D_{max}	WL	lyrs	D_{max}	WL	lyrs	D_{max}	WL	lyrs	D_{max}
n30	0.619	3.5	2.609	0.561	3.5	2.609	0.601	3.5	2.609	0.592	3.5	2.083
n50	0.899	5	2.722	0.799	4.5	2.722	0.843	5	2.721	0.857	4.5	2.079
n100	1.458	7.5	2.666	1.349	7.5	2.666	1.439	7.5	2.503	1.399	7.5	2.036
n200	2.710	16	2.553	2.524	13	2.553	2.582	16	2.553	2.650	16	2.143
n300	4.098	13	3.237	3.781	12.5	3.059	3.956	13	3.037	4.041	12.5	2.837
TIME	20			108			33			52		

TABLE III
TIMING-DRIVEN OPTICAL ROUTING WITH DIFFERENT CRITICALITY LEVELS. WE REPORT THE TOTAL WIRELENGTH, ROUTING LAYER (XY) USAGE, MAXIMUM NET DELAY (NS), AND MAXIMUM VIOLATIONS

	cLevel=0.90				cLevel=0.85				cLevel=0.80			
ckt	WL	lyrs	D_{max}	viol	WL	lyrs	D_{max}	viol	WL	lyrs	D_{max}	viol
n30	0.606	3.5	2.261	0	0.604	3.5	2.134	0	0.592	3.5	2.083	0
n50	0.884	5	2.283	0	0.872	4.5	2.234	0	0.857	4.5	2.079	0
n100	1.442	7.5	2.356	0	1.426	7.5	2.212	0	1.399	7.5	2.036	0
n200	2.677	16	2.283	0	2.650	16	2.143	0	2.598	16	1.987	1
n300	4.041	12.5	2.837	0	3.979	12.5	2.638	4	3.934	12	2.453	6
TIME	27				48				57			

are 13%, 17.6%, and 22.8%, respectively. However, as the criticality level reduces, the violations for larger circuits grow. The number of violations is an indicator of the difficulty in fixing the violations as a post process. A maximum of 23.7% delay improvement without resource violations is achieved with a criticality level of 0.80 for $n50$ and $n100$.

VI. CONCLUSION

Optical waveguides are useful for extremely high propagation speeds, low crosstalk, and transmission of multiple signals simultaneously. However, device costs and signal losses associated with optical devices call for careful consideration during optical routing within SOP. In this paper, we presented the first optical router for 3-D SOP. We developed algorithms to handle the construction of optical waveguides based on performance and congestion objectives, optimum net-to-waveguide mapping using maximum-flow minimum-cost network flow model, and ripping and rerouting of existing nets to relieve congestion and make maximum utilization of waveguides.

REFERENCES

[1] R. Tummala, "SOP: What is it and why? A new microsystem-integration technology paradigm-moore's law for system integration of miniaturized convergent systems of the next decade," *IEEE Trans. Adv. Packag.*, vol. 27, no. 2, pp. 241–249, May 2004.

[2] G. Chang, D. Guidotti, F. Liu, Y. Chang, Z. Huang, V. Sundaram, D. Balaraman, S. Hegde, and R. Tummala, "Chip-to-chip optoelectronics SOP on organic boards or packages," *IEEE Trans. Adv. Packag.*, vol. 27, no. 2, pp. 386–397, May 2004.

[3] S. Lee, T. Lemczyk, and M. Yovanovich, "Analysis of thermal vias in high density interconnect technology," in *Proc. IEEE Semi-Therm Symp.*, 1992, pp. 55–61.

[4] T. Chiang, K. Banerjee, and K. Saraswat, "Effect of via separation and low-k dielectric materials on the thermal characteristics of CU interconnects," in *IEDM Tech. Dig.*, 2000, pp. 261–264.

[5] C. Seo, A. Chatterjee, and N. Jokerst, "Physical design of optoelectronic system-on-a-package: A CAD tool and algorithms," in *Proc. 6th Int. Symp. Qual. Electron. Design*, 2005, pp. 567–572.

[6] C. Seo, A. Chatterjee, S. Cho, and N. Jokerst, "Design and optimization of board-level optical clock distribution network for high-performance optoelectronic system-on-a-packages," in *Proc. 14th ACM Great Lakes Symp. VLSI*, 2004, pp. 292–297.

[7] R. S. Li, "Optimization of thermal via design parameters based on an analytical thermal resistance model," in *6th Intersoc. Therm. Thermomech. Phenom. Electron. Syst.*, 1998, pp. 475–580.

[8] X. Hong, G. Huang, Y. Cai, S. Dong, C. K. Cheng, and J. Gu, "Corner block list: An effective and efficient topological representation of non-slicing floorplan," in *Proc. IEEE Int. Conf. Computer-Aided Design*, 2000, pp. 8–12.

[9] W. Elmore, "The transient response of damped linear networks with particular regard to wideband amplifiers," *J. Appl. Phys.*, pp. 55–63, 1948.

[10] J. Minz, E. Wong, M. Pathak, and S. K. Lim, "Placement and routing for 3-D system-on-package designs," *IEEE Trans. Compon. Packag. Technol.*, vol. 29, no. 3, pp. 644–657, Sep. 2005.

[11] J. Cong, A. B. Kahng, and K.-S. Leung, "Efficient algorithms for the minimum shortest path steiner arborescence problem with applications to VLSI physical design," *IEEE Trans. Computer-Aided Design Intgr. Circuits Syst.*, vol. 17, no. 1, pp. 24–39, Jan. 1998.

[12] L. McMurchie and C. Ebeling, "Pathfinder: A negotiation based performance-driven router for FPGAs," in *Proc. ACM Int. Symp. Field-Programmable Gate Arrays*, 1995, pp. 111–117.

[13] S. K. Rao, P. Sadayappan, F. K. Hwang, and P. W. Shor, "The rectilinear steiner arboresence problem," *Algorithmica*, pp. 277–288, 1992.

[14] J. Cong, L. He, K. Y. Khoo, C. K. Koh, and Z. Pan, "Interconnect design for deep submicron ICs," in *Proc. IEEE Int. Conf. Computer-Aided Design*, 1997, pp. 478–485.

Jacob Rajkumar Minz (S'05) received the B.Tech. degree in computer science and engineering from the Indian Institute of Technology (IIT), Kharagpur, in 2001 and the Ph.D. degree in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, in 2006.

He was with the Advanced VLSI Design Lab, IIT Kharagpur for a year where he was involved in the design of digital chips. Currently, he is working for Synopsys, Sunnyvale, CA. His areas of interest are physical design automation, logic synthesis, and algorithms for electronic CAD.

Somaskanda Thyagaraja (S'07) received the B.S. degree in computer engineering from the Georgia Institute of Technology (Georgia Tech), Atlanta, in 2005.

He is currently employed as a Systems Engineer at Advanced Micro Devices, Austin, TX. He was with the Georgia Tech CAD Group for one year, where he was involved in research on 3-D system-on-packages and routing algorithms.

Sung Kyu Lim (S'94–M'00–SM'05) received the B.S., M.S., and Ph.D. degrees from the Computer Science Department, University of California, Los Angeles (UCLA), in 1994, 1997, and 2000, respectively.

From 2000 to 2001, he was a Post-Doctoral Scholar at UCLA, and a Senior Engineer at Aplus Design Technologies, Inc. He joined the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, in 2001, where he is currently an Associate Professor. He has been on the Advisory Board of the ACM Special Interest Group on Design Automation (SIGDA) since 2003. He served as a Guest Editor for the *ACM Transactions on Design Automation of Electronic Systems*. His research focus is on the physical design automation for 3-D circuits, 3-D system-on-packages, microarchitectural physical planning, and field-programmable analog arrays.

Dr. Lim received the Design Automation Conference (DAC) Graduate Scholarship in 2003, the National Science Foundation Faculty Early Career Development (CAREER) Award in 2006, and the Outstanding Junior Faculty Member Award from the School of Electrical and Computer Engineering, Georgia Institute of Technology, in 2007. He is an Associate Editor of the IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION SYSTEMS. He has served the Technical Program Committee of several ACM and IEEE conferences on electronic design automation.