Power-Optimal Layout of Passive, Single-Hop, Fiber-Optic Interconnections Whose Capacity Increases with the Number of Stations

Yitzhak Birk
Electrical Engineering Department
Technion – Israel Institute of Technology
Haifa 32000, Israel *

Abstract

Passive, single-hop interconnections among N stations, each with several transmitters and receivers, provide a communication path between any two stations using one of the transmitters and receivers of each station. Such interconnections which permit a polylogarithmic (in N) number of concurrent transmissions using only spatial separation have recently been described. This paper shows how to optimally compose the layouts of any two fiber-optic single-hop interconnections to form that of the "product" interconnection. It presents a general method for optimally transforming a layered directed acyclic graph so as to balance the indegree and outdegree of intermediate vertices without altering (source, destination) connectivity.

1. Introduction

1.1 Single-hop interconnections

A *single-hop interconnection (SHI)* is a static, passive interconnection network that provides a direct, though possibly shared, path among all stations. There are no routing switches, and no forwarding by intermediate stations is required. In its most basic form, which is assumed in this paper, the use of tunable components or code-division multiple access is disallowed, thus requiring only the simplest and least expensive components. Such an SHI among single-transmitter and -receiver stations permits only one ongoing transmission. Examples include Ethernet [1], computer buses and single-frequency local radio networks.

Equipping each station with several transmitters and receivers permits the construction of a variety of SHIs, which fall into two categories [2]:

- *Bus-oriented* SHIs. These can be described as a collection of buses, with each transmitter and receiver connected to exactly one bus. The sets of receivers that can hear any two transmitters are either identical or disjoint.
- *Non-bus-oriented* SHIs. There are at least two transmitters such that the sets of receivers that can hear them are neither identical nor disjoint.

For clarity of exposition, it is convenient to think of an SHI as connecting a set of m source stations (SSs), each with c_T transmitters, to a set of n destination stations (DSs), each with c_R receivers. Whenever $m = n$, (SS,DS) pairs can be used to form bidirectional stations.

1.2 Bus-oriented SHIs

As depicted in Fig. 1, it is possible to construct up to $c_T \cdot c_R$ equally-populated buses such that any two stations have at least one bus in common [2][3]. Skewed traffic patterns may result in substantially lower concurrency than $c_T \cdot c_R$. When implemented in fiber-optic technology, this bus-oriented interconnection is also optimal in terms of path loss [4].

1.3 Non-bus-oriented SHIs

Consider a bipartite graph with "transmitter" vertices on the left and "receiver" vertices on the right.

*Part of this research was carried out at IBM's Almaden Research Center. Email: birk@ee.technion.ac.il

0743-166X/93 $03.00

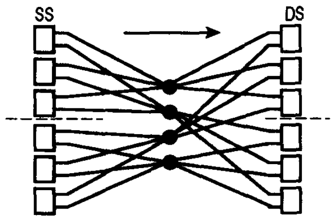

Figure 1: A single-path bus-oriented SHI. $c_T = c_R = 2$.

It defines a *Shared Directional Multichannel* (SDM) whose inputs and outputs correspond to the "transmitter" and "receiver" vertices, respectively. A signal presented at a channel input reaches all the outputs that are connected to this input by graph edges [2][5]. A message is received successfully by a receiver connected to an output of an SDM if and only if it is addressed to that receiver and the receiver hears no other transmissions. An SDM needn't be bus-oriented.

An SDM can be used as an SHI among stations with multiple transmitters and receivers. We use $(c_T, c_R; m, n)$ to denote the size of such an interconnection. A $(2,1;N,N)$ interconnection with uniform-traffic concurrency $\log_2 N$ at high loads was described in [5]. In [6], the concurrency was increased by one. SDM-based SHIs whose uniform-traffic concurrency at heavy loads increases with N as $(\log_2 N)^{c_T+c_R-2}$ when operated with a specific round-robin transmission schedule were also presented in [5].

1.4 Fiber-optic SHIs

A fiber-optic SHI consists of transmissive star couplers interconnected via fiber segments. Unlike with a "star" of copper wires, a signal presented at an input of a transmissive fiber-optic star coupler is split among all outputs, but does not return over the input lines.

In an $(x \times y)$ lossless star coupler with equal coupling to all outputs, the relation between input power, P_{in}, and the power at each output, P_{out}, is

$$\frac{P_{in}}{P_{out}} = \max(x, y).$$

This somewhat unintuitive fact is sometimes referred to as the *"fan-in"* problem in fiber optics [7][8].

In fiber-optic implementations with direct detection, the maximum permissible transmission rate is inversely proportional to the power loss along the path (*"path loss"*) [9][10], corresponding to a requirement for a minimal amount of energy per bit at the receiver. In view of this, the fan-in problem and the fact that capacity is the product of transmission rate and concurrency, it is important to take layout into account when comparing the capacities of fiber-optic SHIs. Specifically, naive layouts of non-bus-oriented SHIs are impractical. For facility of exposition, we will assume lossless components, so path loss is only due to splitting and merging of signal paths.

1.5 The layout challenge

The layout-optimization challenge can be stated as follows: replace the bipartite graph representing a given SHI with a graph that may have intermediate vertices (corresponding to star couplers) such that the maximum over (transmitter, receiver) pairs of the product of max(indegree, outdegree) of the vertices along the path is minimized. Transmitter and receiver vertices are included in the paths. The new graph must nevertheless provide a path between a transmitter and a receiver if and only if there was an edge between them in the bipartite graph. We are also interested in complexity, whose measures are the numbers of fiber segments or star couplers.

In an SHI of size $(c_T, c_R; m, n)$, there must be a transmitter that reaches at least n/c_T receivers and a receiver that can hear at least m/c_R transmitters. Therefore, a lower bound on path loss for any such SHI is $\max(n/c_T, m/c_R)$. The maximum-capacity bus-oriented SHIs can easily be laid out optimally in this respect and also have minimum complexity, requiring $N \cdot c/2 \cdot log_2 N$ couplers of size (2×2) for $c_T = c_R = c$ and $m = n = N$. (The number of fiber segments is approximately twice as large as the number of (2×2) couplers.) If there is no restriction on coupler size, the number of fiber segments is $(m \cdot c_T + n \cdot c_R)$.[4]

The naive layout of an SDM-based SHI as a bipartite graph requires a $(1 \times \frac{n}{c_T})$ coupler connected to each transmitter and an $(\frac{m}{c_R} \times 1)$ coupler connected to each receiver. In the common case of $m = n = N$, the path loss would be $N^2/(c_T \cdot c_R)$ and approximately N^2/c elementary couplers would be required. As pointed out in [6], the concurrency advantage of the SDM-based SHIs over the bus-oriented ones could easily be offset by the need for slower transmissions.

In [11], a power-optimal layout for a $(2,1;k,2^k)$ SDM-based SHI with concurrency $\log_2 N$ was presented. In this paper, we present a method for composing the layouts of two existing interconnections into one for the "product" interconnection, each of whose parameters is equal to the product of the respective values in the two given interconnections, while preserving the quality of the layout. This complements the method presented in [5] for composing large SHIs from small ones in terms of connectivity and schedule. The composition method is based in part on a novel method for improving the balance between indegree and outdegree of vertices in a layered directed acyclic graph, and is the main contribution of this paper.

Section 2 briefly reviews SDM-based SHIs. Section 3 then presents the layout-composition method and provides examples. Wiring complexity is addressed in section 4, and section 5 offers concluding remarks.

2. High-concurrency SHIs

An SHI that provides a single path between each SS and every DS can be specified by a *wiring* $W(s,d) = (W_1(s,d), W_2(s,d))$, with (s,d) denoting an (SS,DS) pair. $W_1(s,d)$ specifies the transmitter used by s to reach d and $W_2(s,d)$ — the receiver used by d to listen to s. Whenever an SHI is operated with a fixed round-robin schedule, we use $X(s,d)$ to specify the time slot in which s may transmit its message to d. (We assume a fixed message-transmission time of one time slot.)

2.1 A $(2,1;N,N)$ interconnection with concurrency $\log_2 N$ [5]

Consider initially an SHI of size $(2,1;k,2^k)$. Here, $W_1(s,d) \in \{0,1\}$ and $W_2(s,d) \equiv 0$. In [5], it was shown that by numbering the DSs with k-bit numbers and having the ith SS select which of its transmitters to use based on the ith bit in the DS's address, it is possible to construct a fixed round-robin transmission schedule that permits a single message to be successfully sent from each SS to every DS in 2^k time slots, i.e., k concurrent transmissions. The interconnection is depicted in Fig. 2, using only one "layer" of SSs.

In [6], this was improved upon by presenting a $(2,1;k+1,2^k)$ SHI with concurrency $k+1$. The connections and schedule for the first k SSs are unchanged; the last SS picks a transmitter based on the parity of the bit string constituting the DS address.

As shown in Fig. 2, the number of SSs in either of these SHIs can be increased to equal the number of DSs by adding groups of k (respectively $k+1$) SSs.

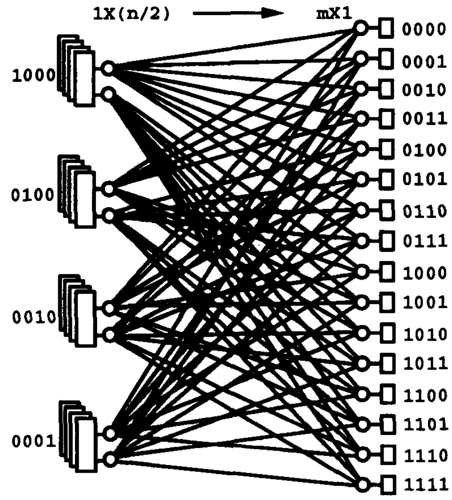

Figure 2: The logical interconnection and naive layout of an interconnection of 2^k 2-transmitter SSs and 2^k single-receiver DSs. ($k=4$.) The SSs are of k "types", each wired differently, with $2^k/k$ identically-wired SSs in each type. Rectangles and circles denote stations and star couplers, respectively.

SSs in the same position in different groups all have identical wiring. We say that they are of the same *type*. The concurrency can be retained by repeating the original schedule N/k times, once for each group of SSs. (Divisibility problems have a minor impact and are ignored for brevity.) This yields a $(2,1;N,N)$ SHI with concurrency $\log_2 N$ for uniform traffic and high load.

Efficient layout of $(2,1;k,2^k)$ [11].

The only way to meet the implementation challenge is to overlap the splitting and merging of signal paths, so as to avoid unnecessary excess fan-in. Stated differently, we must strive to balance the couplers, i.e., equate the number of inputs and outputs of each coupler. This was achieved in [11] by taking advantage of wiring symmetries among SSs and among DSs. The construction in [11] produced three stages of couplers. Coupler sizes in the respective stages are

$$(1 \times 2^{k/2-1}),\ \ (\frac{k}{2} \times 2^{k/2}),\ \text{and}\ (2 \times 1), \qquad (1)$$

yielding a path loss of 2^k. This appears to be sub-optimal by a factor of two. However, the excess fan-in in the last stage is unavoidable since no two receivers hear identical sets of transmitters, so the layout is power-optimal. Fig. 3 depicts the optimal layout for $k = 4$. A comparison with Fig. 2 reveals the savings, which become much more pronounced as N increases.

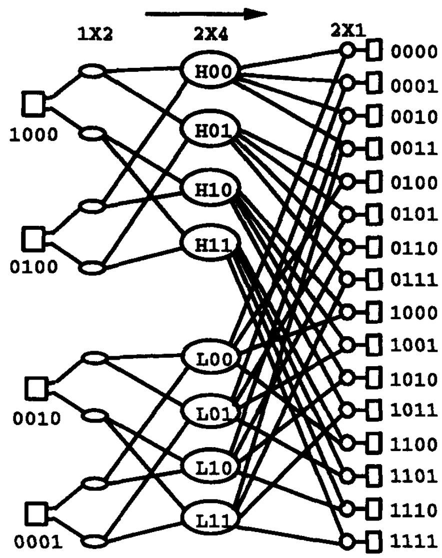

Figure 3: Power-optimal layout of an interconnection of k 2-transmitter SSs and 2^k single-receiver DSs. ($k = 4$.) All couplers are balanced, except for the unavoidable (2×1) stage-3 couplers. The path loss is 16, which is optimal.

Lemma 1. *A power-optimal layout for an SHI of size* $(c_T, c_R; m, n)$*, when reversed, is optimal for the corresponding SHI of size* $(c_R, c_T; n, m)$.

Proof. By contradiction. If the reversed layout is not optimal, then (given the faithful representation of the wiring) there is excess fan-in in one of its stages and excess fan-out in another, such that the two could actually be combined in the same stage to avoid the excess fan-in. Clearly, the same is true of the original layout with the roles of fan-in and fan-out reversed. □

Power-optimal layouts for the $(2, 1; k+1, 2^k)$ and for $(2, 1; 2^k, 2^k)$ SHIs were also constructed in [11]. The power-optimal layout for the latter will be derived here in a much easier way as a simple application of the new composition method.

2.2 Composing large interconnections from smaller ones

In [5] it was shown how, given two SHIs, to construct a "product" SHI: the value of each of its parameters (m, n, c_T, c_R and the concurrency) is the product of the respective parameter values in the constituent SHIs. The construction, which will be used in this paper to verify the correctness of the layouts, is as follows.

Let $i \in \{1, 2\}$ index two given SHIs. For $i \in \{1, 2\}$, let S_i, D_i, T_i, R_i denote the sets of SSs, DSs, transmitters of each SS and receivers of each DS, respectively. Thus, $|S_i| = m_i, |D_i| = n_i, |T_i| = c_{T_i}$, and $|R_i| = c_{R_i}$.

Let $W^i = (W_1^i, W_2^i)$ and X^i be a compatible wiring and schedule for the communication between S_i and D_i, such that $W_1^i(s_i, d_i) \in T_i$ and $W_2^i(s_i, d_i) \in R_i$ for $(s_i, d_i) \in S_i \times D_i$ (Cartesian product).

Let $S = S_1 \times S_2$, $D = D_1 \times D_2$, $T = T_1 \times T_2$ and $R = R_1 \times R_2$. For a pair of stations ($s = (s_1 s_2) \in S$, $d = (d_1 d_2) \in D$):

$$\begin{aligned} W(s,d) &= (W_1(s,d), W_2(s,d)) \qquad (2) \\ &= ((W_1^1(s_1,d_1) W_1^2(s_2,d_2)), (W_2^1(s_1,d_1) W_2^2(s_2,d_2))) \end{aligned}$$

and $X(s,d) = (X^1(s_1,d_1) X^2(s_2,d_2))$.

Note. "$(x_1 x_2)$" denotes the concatenation of the two numbers x_1 and x_2 to form a new number. If x_1 and x_2 are in different bases, the concatenation results in a "hybrid" number, whose range of values is nevertheless $[0, \max(x_1) \cdot \max(x_2)]$.

3. Efficient composition of layouts

In this section we present the main contribution of this paper: a composition method for layouts which complements the method of composing wiring functions presented in [5] and summarized in Eq. (2).

Given layouts of the two interconnections, #1 and #2, that are to be combined to form a larger one, the construction is carried out in two steps: (i) abutting and (ii) compaction.

3.1 Abutting

The outputs of a column of $m_1 \cdot c_{T_1}$ #2 interconnections are connected to the inputs of a column of $n_2 \cdot c_{R_2}$ #1 interconnections. (The order can be reversed.) Thus, a signal goes through a #2 interconnection followed by a #1 interconnection. The connections are made as follows.

Initially, label each port of a #2 interconnection with the numbers of the station and transmitter (or receiver) connected to it in a stand-alone #2 interconnection. Input and output labels will assume the form $(s_2 \in S_2, t_2 \in T2)$ and $(d_2 \in D_2, r_2 \in R_2)$, respectively.

Next, uniquely label each of the $n_2 \cdot c_{R_2}$ #2 interconnections with a pair of numbers $(s_1 \in S_1, t_1 \in T_1)$. There are exactly enough such pairs, so each input of the first column is uniquely labeled $(s_1 \in S_1, t_1 \in T_1, s_2 \in S_2, t_2 \in T_2)$, and each of its outputs — $(s_1 \in S_1, t_1 \in T_1, d_2 \in D_2, r_2 \in R_2)$. Finally, connect transmitter t_1t_2 of SS s_1s_2 to the input labeled (s_1, t_1, s_2, t_2).

Similarly, uniquely label each of the $n_2 \cdot c_{R_2}$ #1 interconnections $(d_2 \in D_2, r_2 \in R_2)$; uniquely label each output of the second column $(d_2, r_2, d_1 \in D_1, r_1 \in R_1)$ and connect it to receiver number r_1r_2 of DS number d_1d_2, and label each input of the second column (s_1, t_1, d_2, r_2).

Finally, connect each output of the first column to the input of the second column that has the same label.

Proposition 2. *The foregoing construction faithfully represents the wiring called for by [5].*

Proof. Without loss of generality, let us consider source station number $s_1^*s_2^*$ and destination station $d_1^*d_2^*$. From (2) it follows that the transmitter and receiver used for connecting those two should be $t_1^*t_2^*$ and $r_1^*r_2^*$, respectively, where (t_i^*, r_i^*), $i \in \{1,2\}$, is the (transmitter-number, receiver-number) pair used in a stand-alone $\#i$ interconnection for communication between s_i and d_i.

In the foregoing construction, transmitter $t_1^*t_2^*$ of source-station $s_1^*s_2^*$ is connected to the #2 interconnection labeled (s_1^*, t_1^*). Since the first part of the path is a correctly-constructed #2 interconnection, the outputs of this interconnection which are reached by this transmitter include (d_2^*, r_2^*). Combined with the interconnection's label, this output of the first column is labeled $(s_1^*, t_1^*, d_2^*, r_2^*)$.

Next, let us turn our attention to the second part of the path, namely the #1 interconnection. Receiver $r_1^*r_2^*$ of destination station $d_1^*d_2^*$ is connected to the #1 interconnection labeled (d_2^*, r_2^*). Since this is a correctly-implemented #1 interconnection, the inputs of this interconnection from which this receiver can be reached include (s_1^*, t_1^*). Combined with the interconnection's label, this input is labeled $(s_1^*, t_1^*, d_2^*, r_2^*)$. But this input of the 2nd column is connected to the $(s_1^*, t_1^*, d_2^*, r_2^*)$ output of the first column, so there is indeed a path between the appropriate transmitter and receiver of s^* and d^*.

We must still show that there are no paths other than the required ones. This follows from the fact that each input of a #2 interconnection reaches n_2/c_{T_2} of its outputs, and each input of a #1 interconnection reaches exactly n_1/c_{T_1} of its outputs. Therefore, concatenating the two colums of interconnections makes it possible for each transmitter to reach at most $n_1n_2/(c_{T_1}c_{T_2})$ receivers, which is exactly the number called for by the wiring matrix of (2). Since all the required connections exist, it follows that there are no extra ones. Finally, since only the desired connections exist and each constituent interconnection provides at most one path between each of its inputs and outputs, there is at most one signal path between each input and output of the composite interconnection. □

3.2 Compaction

In this step, we attempt to balance the couplers so that their indegree is as close as possible to their outdegree. The following two lemmas, illustrated in Fig. 4, present the transformations that may be used. (Variable names in the lemmas bear no relationship to names used earlier.)

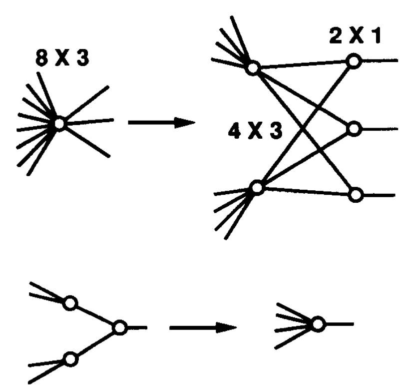

Figure 4: Legal coupler-degree transformations.

Lemma 3. *In terms of connectivity, an $(m \times n)$ star coupler is equivalent to k couplers of size $(m_i \times n)$ with $\sum_{i=1}^{k} m_i = m$, feeding n couplers of size $(k \times 1)$, connected as follows: the inputs to the original coupler are partitioned among the inputs of the $(m_i \times n)$ couplers; each $(k \times 1)$ coupler receives one output of every $(m_i \times n)$ coupler.* □

Lemma 4. *An $(n \times 1)$ coupler whose output feeds another coupler can be eliminated by adding $n-1$ inputs to the latter and feeding the inputs of the $(n \times 1)$ coupler directly to it.* □

The indegree of a coupler can thus be reduced in exchange for an appropriate increase to the indegree of couplers in the *next* stage. (The number of couplers in the early stage must also be increased.) Similarly, it is possible to reduce the outdegree of couplers in a given stage by appropriately increasing the outdegree of those in the *previous* stage. The actual changes to connections and coupler structures are readily deduced from the transformation and Fig. 4. Transformations in the opposite directions are not allowed!

One may, for example, reduce the indegree of couplers in one stage by some integer factor k and increase that of couplers in the next stage by a factor of k. One may also swap indegrees of couplers in successive stages, provided that this results in a reduction of indegree to the earlier stage. Similarly, one may swap outdegrees of couplers in successive stages, provided that this results in an increase of the outdegree of the couplers in the earlier stage. Finally, we note that the stages involved in a transformation needn't even be adjacent, provided that the indegree (outdegree) of the earlier stage is reduced (increased). This is so because there is always an equivalent sequence of legal transformations involving degrees of adjacent stages.

Compaction consists of successively applying these transformations with appropriately chosen factors until maximum balance is attained. The associated rearrangement of connections can trivially be deduced from the transformations.

3.3 Examples

In the first example, we compose the efficient layout of an SHI of size $(1,1;2^k/k,1)$ with that of $(2,1;k,2^k)$ to produce an efficient layout for the more interesting $(2,1;2^k,2^k)$ SHI, reproducing the result of [11]. Next, we provide a numerical example in which the numbers do not divide well. Finally, we compose the efficient layout of $(2,1;k,2^k)$ with that of $(1,2;2^k,k)$ to produce an efficient layout of $(2,2;k2^k,k2^k)$.

Example 1: $(2,1;k,2^k)+(1,1;\frac{2^k}{k},1) \rightarrow (2,1;2^k,2^k)$. The coupler sizes for an optimal layout of $(2,1;k,2^k)$ are $(1 \times 2^{\frac{k}{2}-1})$, $(\frac{k}{2} \times 2^{\frac{k}{2}})$, (2×1). An SHI of size $(1,1;\frac{2^k}{k},1)$ is trivially implemented using a $(\frac{2^k}{k} \times 1)$ coupler.

Noting that the first interconnection has excess fanout while the second one has excess fan-in, we place the second one in the first part of the path. This yields

$$(\frac{2^k}{k} \times 1),(1 \times 2^{\frac{k}{2}-1}),\ (\frac{k}{2} \times 2^{\frac{k}{2}}),\ (2 \times 1).$$

Combining the couplers of the first and second stages, and then reducing the indegree of stage-1 couplers by a factor of $2^{\frac{k}{2}}/(\frac{k}{2})$ and increasing the indegree of stage-2 couplers by the same factor yields

$$(2^{\frac{k}{2}-1} \times 2^{\frac{k}{2}-1}),\ (2^{\frac{k}{2}} \times 2^{\frac{k}{2}}),\ (2 \times 1).$$

The resulting path loss is 2^k, which is optimal. Fig. 5 depicts the optimal layout for $k = 4$.

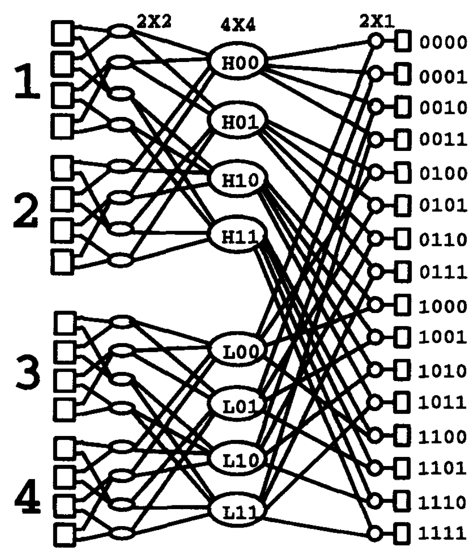

Figure 5: Power-optimal layout of a $(2,1;2^4,2^4)$ SHI. Path loss = 16.

Coupler sizes following each transformation, as well as initial and final sizes, must must all be integer. When a problem occurs, it can be solved by augmentation as illustrated in the following numerical example.

Example 2: $(2,1;6,2^6)$+$(1,1;\frac{2^6}{6},1) \rightarrow (2,1;2^6,2^6)$. The coupler sizes for the constituent interconnections are $(\frac{64}{6}\times 1)$ and (1×4) (3×8) (2×1). Clearly, we have a problem with the first interconnection, which we can only solve by slightly increasing its size. We could increase the number of its SSs to 11 and use an (11×1) coupler. However, we immediately observe that this would prevent any transformations, so we increase the number to 12 and use a (12×1) coupler. The number of SSs in the composite interconnection will be 72, eight of which are dummy. We now have two reasonable transformations: dividing the indegree of the first stage by 3 or by 4. The better one is a division by 3, producing (after composition and compaction):

$$(4\times 4),\ (9\times 8),\ (2\times 1).$$

The path loss is 72, which is optimal for 72 SSs. (Path loss with the other option is 96.)

Example 3: $(2,1;k,2^k)$+$(1,2;2^k,k) \rightarrow (2,2;k2^k,k2^k)$ The coupler sizes for an optimal layout of $(2,1;k,2^k)$, per (1), are:

$$(1\times 2^{\frac{k}{2}-1}),(\frac{k}{2}\times 2^{\frac{k}{2}}),(2\times 1).$$

By symmetry and Lemma 1, the optimal sizes for $(1,2;2^k,k)$ are:

$$(1\times 2),(2^{\frac{k}{2}}\times\frac{k}{2}),(2^{\frac{k}{2}-1}\times 1).$$

Abutting produces an interconnection whose coupler sizes in the various stages are:

$$(1\times 2),(2^{\frac{k}{2}}\times\frac{k}{2}),(2^{\frac{k}{2}-1}\times 1),(1\times 2^{\frac{k}{2}-1}),(\frac{k}{2}\times 2^{\frac{k}{2}}),(2\times 1).$$

Without compaction, the path loss in this interconnection would be 2^{2k}, i.e., approximatly $(N/\log_2 n)^2$.

Merging the last coupler in the first half of the path with the first one in the second yields

$$(1\times 2),(2^{\frac{k}{2}}\times\frac{k}{2}),(2^{\frac{k}{2}-1}\times 2^{\frac{k}{2}-1}),(\frac{k}{2}\times 2^{\frac{k}{2}}),(2\times 1).$$

Next, we reduce the indegree of stage-3 couplers and increase that of stage-4 couplers by a factor of $(2^{\frac{k}{2}}/k)$. Similarly, we reduce the outdegree of stage-3 couplers and increase the outdegree of stage-2 couplers by this factor. The result:

$$(1\times 2),(2^{\frac{k}{2}}\times 2^{\frac{k}{2}-1}),(\frac{k}{2}\times\frac{k}{2}),(2^{\frac{k}{2}-1}\times 2^{\frac{k}{2}}),(2\times 1).$$

Finally, we reduce the indegree of stage-2 couplers and outdegree of stage-4 couplers by a factor of 2 and double the indegree and outdegree of stage-3 couplers, yielding:

$$(1\times 2),(2^{\frac{k}{2}-1}\times 2^{\frac{k}{2}-1}),(k\times k),(2^{\frac{k}{2}-1}\times 2^{\frac{k}{2}-1}),(2\times 1).$$

The path loss here is $k\cdot 2^k = N$. This is optimal for $(2,2;k2^k,k2^k)$, since the excess fan-in in the last stage is again unavoidable.

In the examples, we only manipulated the indegrees and outdegrees of couplers; we did not calculate the number of couplers in each stage and did not specify the actual connections that need to be made. However, these are obvious from the transformations and Fig. 4.

4. Wiring complexity

One measure of complexity is the required numbers of wire segments or fibers. Another is the required number of elementary star couplers [12]. The number of stages traversed by a signal is of lesser interest, as these introduce a very small delay, but will nevertheless be addressed. The comparison will be made with a single broadcast bus interconnecting N stations, c broadcast buses each interconnecting all N stations, and the highest-concurrency bus-oriented SHI, which comprises c^2 buses, each connecting N/c SSs to N/c DSs.

Single bus. With unbounded coupler sizes, a single $(N\times N)$ star coupler and $2N$ fiber segments are required. If (2×2) couplers are used, the interconnection uses $\frac{N}{2}\log_2 N$ star couplers and $N(\log_N +1)$ fiber segments. The signal goes through a single stage and $\log_2 N$ stages in the two cases, respectively.

c buses, each interconnecting N **stations.** These require c times more components than for the single bus.

Maximum-concurrency Bus-oriented SHI. Assuming unbounded coupler sizes, this requires c^2 couplers of size $(\frac{N}{c}\times\frac{N}{c})$ and $2Nc$ fiber segments. A signal travels through a single coupler. With (2×2) couplers, the number of couplers is $c^2\cdot(\frac{N}{2c}\cdot\log_2\frac{N}{c})$, i.e., $c\cdot(\frac{N}{2}\cdot\log_2\frac{N}{c})$. This is somewhat less than for c broadcast buses, as is the corresponding number of fiber segments. A path comprises $\log_2\frac{N}{c}$ stages.

SDM-based high-concurrency SHIs. With unbounded coupler sizes, the number of coupler stages for $c = 2$ is 5, and it increases only as $\log c$, even without compaction. The number of fiber segments is approximately $N \cdot c$ times the number of stages. With (2×2) couplers, the number of couplers is at most $\frac{N \cdot c}{2} \cdot \log_2 N$, i.e., no larger than the c broadcast buses. The number of stages in this case is $\log_2 N$.

The complexity of the SDM-based SHIs is thus similar to that of the simpler bus-oriented SHIs which offer substantially lower concurrency.

5. Summary

Non-bus-oriented single-hop interconnections based on a shared directional multichannel can offer much higher uniform-traffic concurrency than bus-oriented ones. Moreover, they can be implemented efficiently both in terms of path loss and fiber-optic component requirements. Efficient layouts are derived directly for small interconnections and then, using the composition method presented in this paper, are composed to form efficient layouts for much larger ones. Ignoring integer constraints, the general layout method presented in this paper produces implementations with a path loss (with lossless components) of at most N, the number of stations. This is nearly as good as the best possible layouts of power-optimal bus-oriented SHIs.

The significance of the efficient layouts is that the higher concurrency offered by the SDM-based SHIs relative to the bus-oriented ones is not offset by a slower transmission rate.

Present shortcomings of the SDM-based SHIs are the high low-load delay due to the long schedule round, and the small fraction of network capacity that can be allocated to any (source, destination) pair. Presently, work is under way to tackle these problems, and some improvements have already been made.

Finally, a note on technology. While fiber-optics is a natural technology for the SDM-based SHIs due to the directionality of star couplers, diodes or logic gates (AND, OR) could be used in place of directional star couplers to construct SDMs with copper wires.

References

[1] R.M. Metcalfe and D.R. Boggs, "Ethernet: distributed packet switching for local computer networks", *CACM*, vol. 19, no. 7, pp. 395-403, 1976.

[2] Y. Birk, Concurrent Communication among Multi-Transceiver Stations via Shared Media, Ph.D Diss., Elec. Engr. Dept., Stanford Univ., Dec. 1986. Also tech. rep. CSL-TR-87-321, March 1987.

[3] Y. Birk, F.A. Tobagi and M.E. Marhic, "Bus-oriented interconnection topologies for single-hop communication among multi-transceiver stations", *Proc. IEEE INFOCOM '88*, Mar. 1988.

[4] Y. Birk, "Fiber-Optic Bus-oriented Single-Hop Interconnections among Multi-Transceiver Stations," *IEEE J. Lightwave Tech.*, vol. 9, no. 12, pp. 1657–1664, Dec. 1991

[5] Y. Birk, N. Linial and R. Meshulam, "On the Uniform-Traffic Capacity of Single-Hop Interconnections Employing Shared Directional Multichannels", *IEEE Trans. IT*, vol. 39, no. 1, Jan. 1993. Also IBM Res. Rep. RJ 7859 (72519), Dec. '90.

[6] M.T. Busche, B. Hajek, "On optical interconnection of stations having multiple transmitters and receivers," *Proc. 1990 Int. Symp. on Info. Theo. and its Applications (ISITA '90)*, Hawaii, pp. 967–970, Nov. 1990.

[7] J.W. Goodman, "Fan-in and fan-out with optical interconnections," *Optica Acta*, vol. 32, no. 12, pp. 1489-1496, 1985.

[8] R.W. Boyd, *Radiometry and the Detection of Optical Radiation*, John Wiley & sons, N.Y., 1983.

[9] T.V. Muoi, "Receiver design for high-speed optical-fiber systems," *IEEE J. Lightwave Tech.*, vol. 2, pp. 243-267, June 1984.

[10] C.A. Bracket, "Dense Wavelength Division Multiplexing Networks: Principles and Applications", *IEEE JSAC*, vol. 8, no. 6, pp. 948–964, Aug. 1990.

[11] Y. Birk, "Efficient layout of a passive, single-hop fiber-optic interconnection (among N stations) with capacity $\log_2 N$", Proc. 29th Annual Allerton Conf. on Commun., Control and Computing, Oct. 1991, Monticello, IL, pp. 322–331.

[12] M.E. Marhic, "Hierarchic and Combinatorial Star Couplers," *Opt. Let.*, vol. 9, p. 368, Aug. 1984.