Efficient Layout of a Passive, Single-Hop, Fiber-Optic Interconnection (among N stations) with Capacity $\log_2 N$*

Yitzhak Birk
IBM Almaden Research Center
San Jose, California †

Abstract

A passive, single-hop, fiber-optic interconnection among n stations, each with two transmitters and one receiver, and a round-robin transmission schedule for it, which permit $\log_2 n$ concurrent, non-interfering transmissions, have recently been described. This is a substantial improvement over the previously known limit of two concurrent transmissions, but the layout of this interconnection poses a challenge both in terms of wiring complexity and power budget. We show how to lay it out using only $n \log_2 n$ fiber segments. Moreover, the ratio of transmitted to received power is n, which is optimal.

1. Introduction

1.1 Single-Hop Interconnections

A single-hop interconnection ($\mathcal{SHI}$) is a static interconnection which provides a communication path between any two stations. Such paths may be shared, but there is no routing or forwarding. Examples of such interconnections are computer buses, Ethernet, and local radio networks. The capacity of an interconnection for a given traffic pattern is the product of the transmission rate and the number of concurrent non-interfering transmissions ("*concurrency*").

Whenever each station is equipped with a single transmitter and receiver, the only possible topology is a single "bus" interconnecting all stations. With baseband transmissions, the concurrency is one. Equipping stations with multiple transmitters and receivers, however, permits the construction of a variety of $\mathcal{SHI}$'s, which fall into two classes [1]:

- *bus-oriented* $\mathcal{SHI}$'s. These can be described as a collection of buses, with each

*Proc. '91 Allerton Conference

†Presently at the Electrical Engineering Department, Technion–Israel Institute of Technology, Haifa 32000, Israel; birk@ee.technion.ac.il

transmitter and receiver connected to exactly one bus. The sets of receivers that can hear transmissions of any two transmitters are therefore either identical (the two transmitters are on the same bus) or disjoint (they are on different buses).

- *non-bus-oriented* $\mathcal{SHI}$'s. Here, there are at least two transmitters such that the sets of receivers that can hear them are neither identical nor disjoint.

Throughout the remainder of the paper, we will speak of m source stations, SS's, connected to n destination stations, DS's. Whenever $m = n$, however, one can think of these as n pairs of bidirectional stations.

1.2 Fiber-optic $\mathcal{SHI}$'s

A fiber-optic $\mathcal{SHI}$ consists of directional star couplers interconnected via fiber segments. A directional star coupler is a passive element with a number of inputs and a number of outputs. A signal presented at an input is split among all outputs (and only those). This is in contrast with a "star" of copper wires, in which there is no sense of direction. The directionality of the fiber-optic medium makes it most suitable for implementation of $\mathcal{SHI}$'s, particularly the non-bus-oriented ones.

In an $(x \times y)$ star coupler with equal coupling to all outputs, one would expect the power at each output, P_{out}, to be approximately P_{in}/y. Unfortunately, however,

$$\frac{P_{in}}{P_{out}} = \max\{x, y\}. \tag{1}$$

This is known as the "*fan-in*" problem in fiber optics [2].

In fiber-optic implementations with direct detection, the maximum permissible transmission rate is inversely proportional to the power loss along the path ("*path loss*") [3],[4]. (This corresponds to a requirement for a minimal amount of energy per bit at the receiver.) In view of this and the fan-in problem, it is important to take path loss into account when comparing the capacities of fiber-optic $\mathcal{SHI}$'s. For facility of exposition, we will assume lossless components, so path loss is only due to signal splitting and merging.

1.3 Bus-oriented $\mathcal{SHI}$'s

Equipping each station with c_T transmitters and c_R receivers permits the construction of up to $c_T \cdot c_R$ buses such that every bus inteconnects the same number of stations and any two stations have at least one bus in common [1]. The concurrency of such an interconnection for a uniform traffic pattern (equal traffic rate between every pair of stations) is $c_T \cdot c_R$, which is the maximum uniform-traffic concurrency of any bus-oriented $\mathcal{SHI}$ [1][5][6]. This interconnection is also optimal in terms of path loss, which is

$$\frac{P_{in}}{P_{out}} = \max\left\{\frac{n}{c_T}, \frac{m}{c_R}\right\}. \tag{2}$$

Fig. 1 depicts such an interconnection with m=n=6 and $c_T = c_R = 2$.

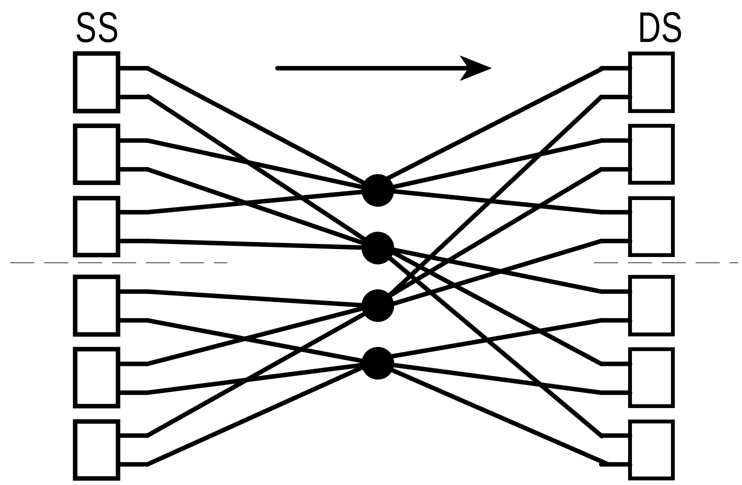

Figure 1: A single-path bus-oriented $\mathcal{SHI}$. There are $c_T \cdot c_R$ buses.

1.4 Shared Directional Multichannels

A shared directional multichannel (SDM, for short) consists of a set of inputs and a set of outputs to which we connect transmitters and receivers, respectively. A signal placed at any given input reaches a subset of the outputs. A channel is specified by the outputs reachable from each input, and it needn't be bus-oriented. (See Fig. 2.) The channel adheres to the following rules [1][7]:

1. A message transmitted into any given input of the channel reaches all outputs connected to it (and the receivers connected to those).
2. A message is received successfully by a receiver iff it is addressed to that receiver and no other signals reach it at the same time.

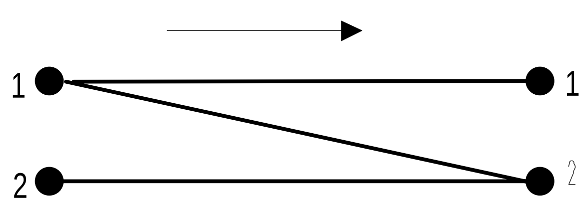

Figure 2: A Shared directional multichannel represented as a bipartite graph. Note that the subsets of receivers that can hear two transmitters are neither identical nor disjoint.

An SDM can also be used to interconnect stations with multiple transmitters and receivers. We use $(c_T, c_R; m, n)$ to denote the size of such an interconnection. For $m{=}n{=}N$, it is possible to construct an SDM whose uniform-traffic concurrency increases with N as $(\log_2 N)^{c_T+c_R-2}$ when operated with a specific round-robin transmission schedule [7]. A $(2,1;k,2^k)$ interconnection with concurrency k was described in [7]. In [6], this interconnection was augmented with another SS, resulting in a $(2,1;k{+}1,2^k)$ interconnection with concurrency $k+1$. Both can be extended to $(2,1;N,N)$ with respective uniform-traffic concurrencies $\log_2 N$ and $1+\log_2 N$.

1.5 The implementation challenge

In this paper, our primary interest is in **uniform-traffic capacity**. However, we are also interested in complexity. The measure of complexity is the number of fiber segments or star couplers. To permit a meaningful comparison, we assume that all couplers are constructed using (2×2) building blocks [8].

Since there must be some transmitter that reaches at least n/c_T receivers and some receiver that can hear at least m/c_R transmitters, a lower bound on path loss for any $(c_T, c_R; m, n)$ $\mathcal{SHI}$ is $\max\{n/c_T, m/c_R\}$. The maximum-capacity bus-oriented $\mathcal{SHI}$'s are optimal 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 optimal layouts for bus-oriented $\mathcal{SHI}$'s occur naturally, but this is not the case for SDM-based $\mathcal{SHI}$'s whose concurrency increases with N. In fact, such an interconnection may require that each transmitter's signal be split n/c_T ways, and that m/c_R signals be combined at the input of each receiver. For the common case of $m{=}n{=}N$, this would require roughly N^2/c elementary couplers, and the path loss would be $N^2/(c_T \cdot c_R)$.

As pointed out in [6], the concurrency advantage of the general SDM-based $\mathcal{SHI}$'s over the bus-oriented ones could be offset by the need for slower transmissions. The challenge is thus to lay out a high-concurrency SDM-based $\mathcal{SHI}$ in a power-efficient manner, so that the concurrency advantage will translate into a capacity advantage. It is moreover desirable to have low complexity.

In this paper, we show how this can be done for the $(2, 1; N, N)$ SDM-based $\mathcal{SHI}$ with concurrency $\log_2 N$ and the improved one of [6]. A generalization of these results to a much broader class of interconnections will be presented elsewhere.

The remainder of the paper is organized a follows. Section 2 describes the $(2, 1; N, N)$ $\mathcal{SHI}$ originally presented in [7] as well as the modification presented in [6]. Section 3 presents an efficient layout for the interconnection of [7], section 4 shows how this can be extended to the interconnection of [6], and section 5 offers some concluding remarks.

2. Specific $(2, 1; N, N)$ Interconnections

2.1 A $(2, 1; k, 2^k)$ interconnection

Let us initially consider k SS's, each with two transmitters, connected to $n = 2^k$ DS's, each with a single receiver. We use k-bit binary vectors to identify the DS's. The k SS's are identified by the k-bit binary vectors that have a single "1".

Wiring. let $W(s, d)$ denote the transmitter ("0" or "1") used by s to reach d. Our interconnection can be described as

$$W(s, d) = s \cdot d. \tag{3}$$

In other words, the ith SS uses the ith bit in the id of each DS to decide which transmitter to connect to that DS. Fig. 3 depicts such an interconnection for $k = 4$.

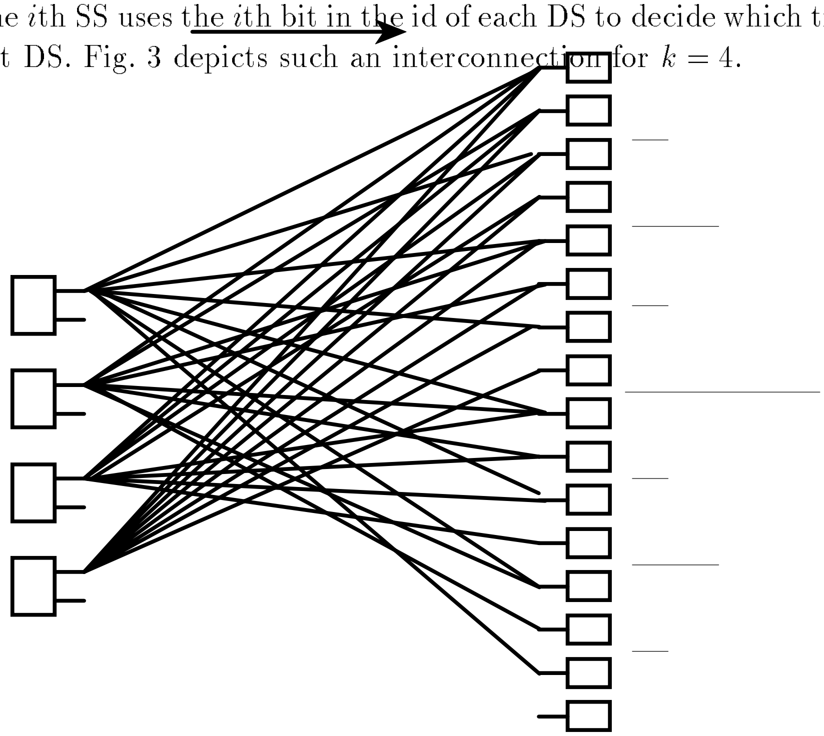

Figure 3: The logical interconnection and straightforward wiring for an interconnection of k 2-transmitter SS's and 2^k single-receiver DS's. ($k = 4$). All missing connections use 2nd-transmitters.

Schedule. $X(s, d)$ specifies the time slot in which s should transmit its message to d using the transmitter $W(s, d)$. A schedule X is *compatible* with a wiring W iff for any two pairs $(s_1, d_1) \neq (s_2, d_2)$ the following holds: if $W(s_1, d_1) = W(s_1, d_2)$, then $X(s_1, d_1) \neq X(s_2, d_2)$. One viable schedule is

$$X(s, d) = s + d. \tag{4}$$

(Modulo 2 arithmetic on the components of the k-bit vectors.) The length of this schedule is 2^k, which is indeed the number of time slots required for transmitting $k \cdot 2^k$ messages, k per slot. See [7] for generalizations and correctness proofs.

The interconnection just described can be extended to $(2, 1; 2^k, 2^k)$ as follows. Partition the SS's into groups of k and apply the foregoing wiring function to the stations within each group. Thus, the ith stations in all groups have identical wiring. Similarly, use the same schedule function to construct a schedule for each group of SS's, and then interleave or concatenate the schedules. The attained concurrency remains k, which is equal to $log_2 N$. We will refer to SS's with identical wiring as being of the same "*type*".

A close inspection of the foregoing interconnections reveals that in any given time slot there is a "free" receiver. This is the basis of the improvement in [6]. The $(2, 1; k+1, 2^k)$ interconnection described there uses the same wiring function for the first k SS's. The last one chooses a transmitter based on the parity of the binary string constituting the id of the receiving station. A schedule is also presented there, but is omitted here for brevity.

3. Efficient Layout of the $(2, 1; N, N)$ Interconnection

The only way to meet the implementation challenge is to overlap the splitting and merging, so that each coupler in the path has equal or nearly equal numbers of inputs and outputs. We achieve this by taking advantage of wiring symmetries among SS's, symmetries among DS's, and overlapping the splitting phase with the merging phase.

Symmetries among SS's. We may merge the outputs of individual transmitters with identical connections (there are $2k$ sets, each with N/k such transmitters) using $(\frac{N}{k} \times \frac{N}{2})$ star couplers. In so doing, we both replicate each of the signals $N/2$ times and combine sets of N/k signals that are to reach the same receivers. For any given receiver, we now take one output fiber of each of the k appropriate couplers and connect them to a $(k \times 1)$ coupler, whose output is connected to the receiver. The total path loss of this scheme is $Nk/2 = N \log_2 N/2$, a substantial improvement over the straightforward scheme. If N/k is not an integer, we add dummy SS's until their number is an integer multiple of k. In other words, we use $\lceil N/k \rceil$ for the group size.

Symmetries among DS's. Recalling that the transmitter used by the ith SS to reach DS number j is determined by the value of the ith bit in the binary representation of j, it follows that all DS's whose numbers (in binary representation) have some x bits in common have the same connections to the corresponding $(x \cdot N/k)$ SS's. The number of such DS's is 2^{k-x}. This leads us to a three-stage layout, which we now begin to construct.

Since no two DS's have identical connections, the signals must reach a DS through a coupler with one output and at least two inputs. Optimistically, we assume that stage-3 couplers are (2×1). This, in turn, implies that an output of a stage-2 coupler must carry signals from all members of some $N/2$ SS's of $k/2$ types. Specifically, we will let such a coupler carry signals from either the first or last $k/2$ types. Since the number of DS's with the first or last $k/2$ bits in common is $2^{k/2}$, this will be the number of outputs of couplers in this stage. Each 3rd-stage coupler will be connected to the outputs of the two stage-2 couplers that carry the combined signals of the transmitters to which its receiver should be connected. (An output of a coupler reached by the appropriate combination of transmitters of the first $k/2$ types of SS's and one carrying the appropriate combination from the remaining $k/2$ types.) The number of couplers in the 2nd stage will be $2 \cdot 2^{k/2}$.

Overlapping the splitting and merging. We have thus far observed that (i) the signals of up to N/k transmitters may be merged, since they all have identical destinations, (ii) a last (3rd) stage of couplers is ideally of size (2×1), and stage-2 couplers each have $2^{k/2}$ outputs. Having determined the number of outputs of couplers in stages 2 and 3, and since each transmitter must reach $2^k/2$ receivers, it follows that the number of outputs of a stage-1 coupler is $(2^{k/2-1})$.

If we took full advantage of the merging possibilities in the first stage, a stage-2 coupler would only have $k/2$ inputs, leading to suboptimal path loss. Instead, we break the merger of N/k signals into two steps: in the first step, we combine groups of x signals (transmitters with identical connections) using $x \times 2^{k/2-1}$ couplers. (x is an integer whose value has yet to be determined.) Next, we take one output of each coupler and connect those outputs to inputs of a stage-2 coupler, which (in addition to other roles) completes the merger. The number of inputs of a stage-2 coupler will therefore be $k \cdot z/2$ for some integer z. Finally, we select the best values of x and z.

Summarizing the situation:

- Stage-3 couplers are (2×1).
- stage-2 couplers have $k \cdot z/2$ inputs and $2^{k/2}$ outputs.
- stage-1 couplers have x inputs and $2^{k/2-1}$ outputs.
- $z \cdot x = N/k = 2^k/k$. (The number of SS's of a given type.)

We now need to pick an (x, z) combination that minimizes the power split. Optimistically, we begin by setting $x = 2^{k/2-1}$. This yields $z = N/(kx) = 2^k/(k \cdot 2^{k/2-1}) = 2^{k/2+1}/k$. Thus, the number of inputs of a stage-2 coupler is $(k/2) \cdot z = 2^{k/2}$. Surprisingly, this is exactly the number of outputs of stage-2 couplers, so we are able to achieve our goal. Fig. 4 depicts the optimal wiring scheme for $k = 6$. Couplers are represented by circles or ellipses, and stations by rectangles.

Final adjustments. The numbers of inputs and outputs of a coupler are all integers, yet there is no guarantee that all the expressions listed above produce integer results. We solve this problem by trying out the nearest integer values and using augmentation in the construction of the interconnection.

3.1 Examples

Example 1: $k = 6$; $N = 2^6 = 64$. (See Fig. 4.)

There are 64 stage-3 couplers, one per DS, each of size (2×1).

There are $2 \cdot 2^{k/2} = 16$ stage-2 couplers. The first eight represent all combinations of transmitter choices from the first $k/2 = 3$ "types" of SS's, and the remaining eight represent the choices from the remaining three types of SS's.

The number of outputs of a stage-2 coupler is $2^{k/2} = 8$.

The number of outputs of a stage-1 coupler is$2^{k/2-1} = 4$.

Since we want $\frac{k}{2} \cdot z = 3z = 8$, the two choices for z are 2 and 3.

Picking $z = 2$, $x = \lceil N/(k \cdot z) \rceil = 6$. The resulting power split is $6 \cdot 8 \cdot 2 = 96$.

Picking $z = 3$, $x = \lceil N/(k \cdot z) \rceil = 4$. The resulting power split is $4 \cdot 9 \cdot 2 = 72$. We therefore pick $z = 3$ and $x = 4$.

Having picked x, we augment the number of SS's to the smallest number which is an integer multiple of $k \cdot x$ and is greater than or equal to N, 72 in this case, and construct the interconnection as follows:

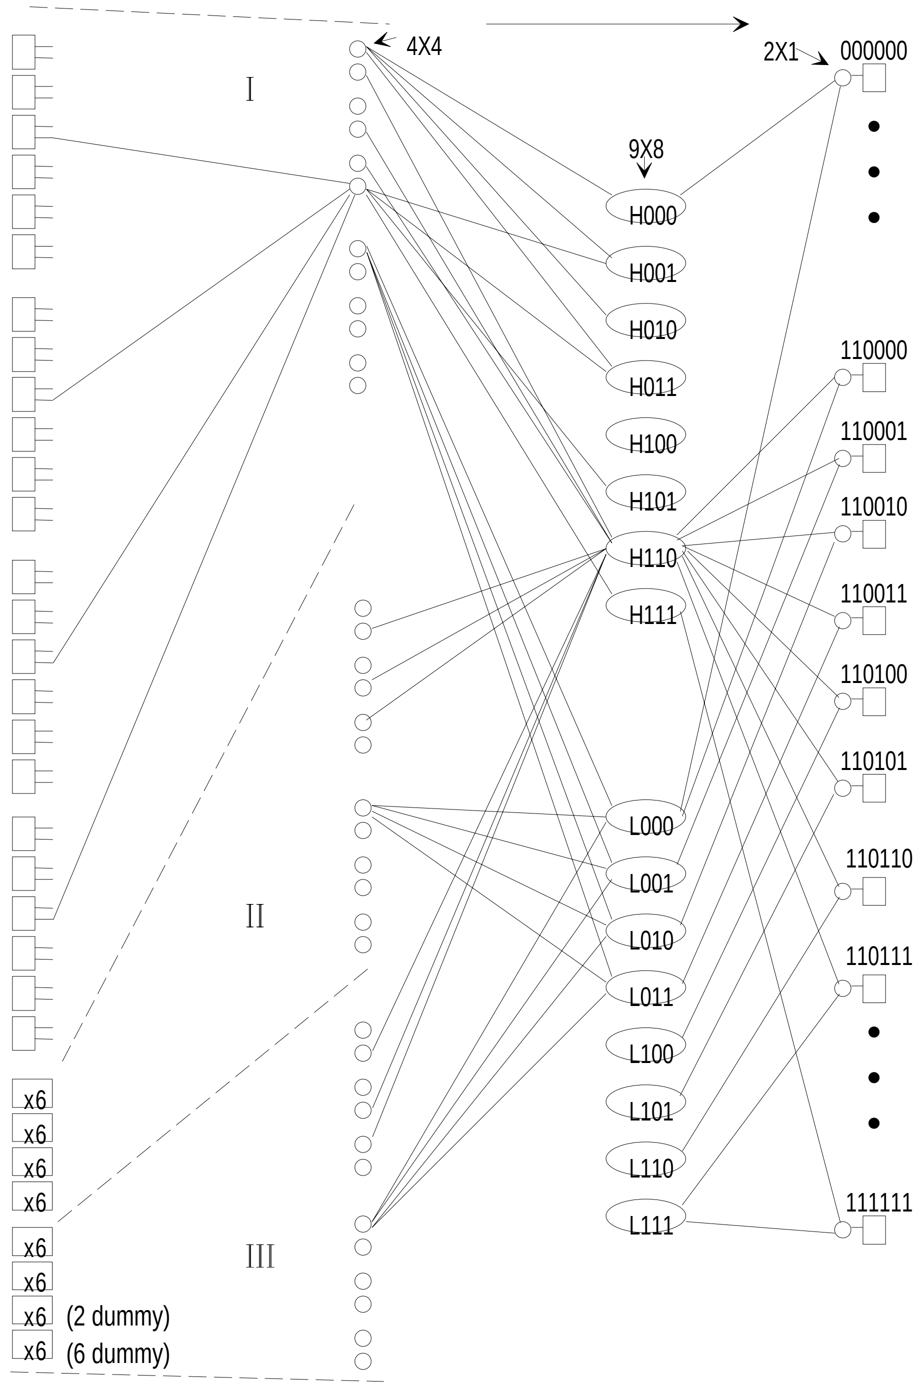

Figure 4: Wiring scheme 2 for $k = 6$; $N = 2^6 = 64$. The power split is 72.

- arrange the SS's in 12 groups of 6 (each contains one SS of each type).
- arrange the groups in 3 clusters of 4 groups.
- in each cluster, there are now 4 SS's with identical connections. Pick each set of 4 transmitters with identical connections and connect them to the four inputs of a 4×4 coupler. We can now think of the outputs of the first stage as three sets of SS's, each of which has one SS of each of the types, and each transmitter of each SS has four output lines carrying identical signals.
- connect outputs of the first stage to inputs of the second one following the example of the 7th stage-2 coupler in Fig. 4, identified as "H110". (The "H" stands for the first half of "types" of SS, i.e., those that have the one in one of the $k/2$ highest-order bit positions.) We connect to it the stage-1 couplers representing the 2nd transmitter of the first SS in each set, the 2nd transmitter of the 2nd SS in each set, and the 1st transmitter of the 3rd SS in each set.
- connect one input of a stage-3 coupler to the output of the stage-2 coupler marked Hxxx, where xxx is the value of the first (high-order, most significant) three bits in the DS number, and the other input to the output of the stage-2 coupler marked Lyyy, where yyy is the value of the three least significant bits in the DS number.

Example 2: $k = 8$; $N = 2^k = 256$.

Here, $2^{k/2-1} = 8$ and $z = 4$. Since z is an integer, this is our solution. Also, no augmentation is required. The power split is $8 \cdot 16 \cdot 2 = 256$, which is perfect.

4. Efficient Layout Based on $(2, 1; k + 1, 2^k)$

The $(2, 1; 2^k, 2^k)$ interconnection derived from the $(2, 1; k + 1, 2^k)$ one described in [6] appears to be more difficult to lay out efficiently, since it is harder to discover symmetries when the wiring function for the SS's of type $(k + 1)$ appears to differ from those for the other types. However, the same wiring rule can be described differently, making the application of our technique to this case straightforward.

We begin by numbering the DS's using the $(k + 1)$-bit numbers with an even number of "1"s. Clearly, there are exactly 2^k such numbers. Moreover, one can readily observe that the last bit of each number is equal to the parity of the string consisting of the previous bits. Finally, the number of stations which have common values in $(k+1)/2$ of their bits is exactly half of that in an interconnection of of the original type with 2^{k+1} stations. To lay out the interconnection, we follow the same recipe as before, making the obvious modifications to accommodate the new numbers.

It is worth noting that this scheme applies most naturally to odd values of k, whereas the original one was most suitable for even ones. Nevertheless, either technique can be used in either case at some penalty.

5. Summary

We have shown how to implement a non-bus-oriented $(2, 1; N, N)$ $\mathcal{SHI}$ with concurrency $\log_2 N$ and a path loss of only N, which is the absolute optimum. While we calculated path loss assuming lossless couplers, the comparison with other interconnections is equally valid for imperfect ones.

Since there is a single path from each SS to each DS and the couplers are all balanced (except for the (2×1) in the last stage), it follows that the required number of elementary couplers and fiber segments is approximately $N \cdot log_2 N$, again optimal to within a constant. Generalizations and extensions of the layout techniques presented here to other values of c_T and c_R, will be presented elsewhere.

In summary, then, we have shown that the shared directional multichannel can offer a significant advantage over buses, permitting the efficient construction of non-bus-oriented $\mathcal{SHI}$'s whose capacity increases with network size.

References

[1] Y. Birk, Concurrent Communication among Multi-Transceiver Stations via Shared Media, Ph.D Dissertation, Electrical Engr. Dept., Stanford University, Dec. 1986. Also available as technical report CSL-TR-87-321, March 1987.

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

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

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

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

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

[7] Y. Birk, N. Linial and R. Meshulam, "On the Uniform-Traffic Capacity of Single-Hop Interconnections Employing Shared Directional Multichannels", IBM Research Report RJ 7859 (72519), Dec. '90.

[8] M.E. Marhic, *Opt. Lett.*, vol. 9, p. 368. 1984.