Power-Efficient Layout of a Fiber-Optic Multistar that Permits $\text{Log}_2\ N$ Concurrent Baseband Transmissions Among N Stations

Yitzhak Birk

***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 jointly permit $\log_2 N$ concurrent noninterfering transmissions on a single wavelength, has 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 path loss. In this paper, a power-efficient implementation of this interconnection using several stages of balanced fiber-optic star couplers is presented. With lossless components, path loss is N, the same as a single-star interconnection that permits only a single transmission at a time. Consequently, the high degree of parallelism translates into higher capacity. The required number of (2×2) star couplers is also very similar to that required for implementing a single $N \times X)$ star.**

***Index Terms*— single-hop interconnections, FOLAN/s, power-efficient layout, shared directional multichannel, multi-star interconnections, parallelism.**

I. Introduction

A. Single-Hop Interconnections

A single-hop interconnection (SHI) is a static interconnection network which provides a communication path between any two stations. Such paths may be shared, but there is no need for routing or forwarding. Examples of such interconnections are computer buses, the Ethernet local-area network [1], and local radio networks. SHI's are attractive because of their simplicity and reliability.

Traditionally, SHI's have been synonymous with a single shared channel, permitting only a single transmission at any given time. Consequently, the peak transmission rate of a station had to exceed the aggregate network throughput. The proliferation of networked PC's, diskless workstations, server-based computing environments and data-intensive applications has resulted in an unacceptable situation: placing all stations on a single shared channel would require very-high-speed communication and very expensive network adapter boards for each station, regardless of its own needs.

Manuscript received July 17, 1992; revised October 22, 1992.
The author is with the Electrical Engineering Department, Technion-Israel Institute of Technology, Haifa 32000, Israel.
IEEE Log Number 9207741.

This problem has most commonly been addressed by partitioning the network into several shared channels interconnected through bridges, and using multiple-hop communication. Recently, however, several approaches have been suggested to overcome this problem without giving up on single-hop connectivity. Most notable is the use of multiple wavelengths on a single physical channel to effectively construct multiple channels. This is an attractive approach, especially with the virtually unlimited bandwidth of optical fibers. However, tunable and wavelength-selective components must be used. Alternatively, spatial or wavelength separation (with nontunable components) can be used if stations are equipped with multiple communication ports (transmitters and receivers). The number of channels that can be constructed is smaller, but the simplest and least expensive components can be used.

The network discussed in this paper follows the latter approach, but does so in a new way that provides much higher concurrency than previously attainable. Moreover, there is no direct parallel to this approach in the wavelength domain. Throughout the paper, we will assume baseband transmissions and a single wavelength, so concurrency (several concurrent non-interfering transmissions) can only be attained through spatial separation among signal paths.

Whenever each station is equipped with a single transmitter and receiver, the only possible topology is a single "bus" interconnecting all stations. The concurrency is one. Equipping stations with multiple transmitters and receivers, however, permits the construction of a variety of SHI's, which fall into two classes [2]:

- bus-oriented SHI's. 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 transmissions of any two transmitters are therefore either identical (the two transmittes are on the same bus) or disjoint (they are on different buses).
- non-bus-oriented 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. Typically, however, $m = n = N$ and these are actually N bidirectional stations.

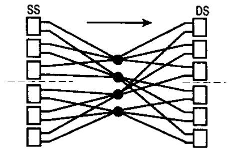

Fig. 1. A bus-oriented SHI. $c_T = c_R = 2$, and there are $c_T \cdot c_R = 4$ buses, each denoted by a large black dot.

B. Fiber-Optic SHI's

A fiber-optic SHI consists of transmissive star couplrs interconnected via fiber segments. Following are several properties of star couplers which are relevant to their use in the implementation of SHI's, as well as some observations on maximum transmission rate.

A signal presented at an input of a transmissive star coupler is split among all outputs, but does not return over the input lines. This is in contrast with a "star" of copper wires, in which there is no sense of direction, and makes fiber-optic technology most suitable for implementation of non-bus-oriented SHI's.

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

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

This somewhat unintuitive fact is sometimes referred to as the "fan-in" problem in fiber optics [3], [4] and will be shown to greatly complicate the fiber-optic implementation of non-bus-oriented SHI's.

In fiber-optic implementations with direct detection, the maximum permissible transmission rate is inversely proportional to the power loss along the path [5], [6]. (This corresponds to a requirement for at least a certain 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 SHI's. For facility of exposition, we will assume lossless components, so path loss is only due to signal splitting and merging.

C. Bus-Oriented 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 interconnects the same number of stations and any two stations have at least one bus in common [2]. 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 SHI [2], [7]–[9]. This interconnection is also optimal in terms of path loss, which is

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

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

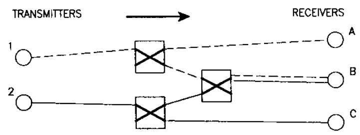

Fig. 2. A Shared directional multichannel constructed using transmissive star couplers: a signal from transmitter 1 reaches receivers A and B, whereas a signal from 2 reaches B and C. This cannot be described as a collection of buses.

D. Non-Bus-Oriented SHI's

As illustrated by Fig. 2, fiber-optic star couplers can be used to construct SHIs in which the sets of receivers that can hear two transmitters are neither identical nor disjoint. This is due to the fact that a signal injected to an input of a coupler is split among the outputs but not the other inputs (unlike the case with copper wires soldered together). An interconnection of this type is referred to as a shared directional multichannel, SDM for short.

A message is received successfully by a receiver if and only if it is addressed to that receiver and no other signals reach the receiver at the same time. (Collisions are only meaningful at individual receivers, as there is no notion of a bus.)

Designs of SDM-based SHI's have been presented in [10]. They consist of wirings and schedules. For each (SS, DS) pair, e.g., (s, d), $W(s, d)$ specifies the transmitter of s and the receiver of d between which the interconnection provides a path, and $X(s, d)$ specifies the time slot (in the round-robin schedule) in which s may send a message to d. We assume single-slot transmissions. 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 [10].

In the remainder of the paper, we restrict the discussion to interconnections between 2-transmiter SS's and single-receiver DS's. (The results also apply directly to the case of a single transmitter per SS and two receivers per DS.) An SHI that permits $\log_2 N$ concurrent transmissions among such stations was described in [10], and in [8] it was improved slightly to permit $\log_2 N + 1$ concurrent transmissions.

E. The Implementation Challenge

Since there must be some transmitter that reaches at least $n/2$ receivers and every receiver must bear m transmitters, a lower bound on path loss for any SHI is $\max\{n/2, m\}$. The maximum-capacity bus-oriented SHI's are optimal in this respect. Our goal is to come as close as possible to this bound.

A straightforward layout of an SDM-based SHI among N stations requires on the order of N^2 fiber segments and (2×2) star couplers. The bus-oriented SHI's, in contrast, are far simpler. Again, we strive to come as close as possible to those.

The optimal layouts for bus-oriented SHI's occur naturally, but this is not the case for SDM-based SHI's whose concurrency increases with N. In fact, the latter may require that each transmitter's signal be split $n/2$ ways at the transmitter output, and that m signals be combined at the input of each

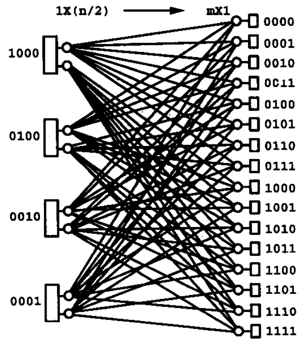

Fig. 3. The logical interconnection and straightforward wiring for an interconnection of k 2-transmitter SS's and 2^k single-receiver DS's. ($k = 4$). Rectangles and circles denote stations and star couplers, respectively.

receiver. For the common case of $m = n = N$, the path loss would be $N^2/2$.

As pointed out in [8], the concurrency advantage of the general SDM-based SHI's over the bus-oriented ones could thus be offset by the need for slower transmissions. The challenge is to lay out a high-concurrency SDM-based 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. Needless to say, the layout must faithfully represent the wiring: the signal of any transmitter must reach exactly those receivers specified by the wiring. In this paper, we show how this can be done. When comparing coupler requirements for different SHI's, we will assume that all couplers are constructed using (2×2) building blocks [11].

The remainder of the paper is organized as follows. Section II describes the SHI originally presented in [10] as well as the modification presented in [8]. Section III presents an efficient layout for the interconnection of [10], Section IV shows how this can be extended to the interconnection of [8], and Section V offers some concluding remarks.

II. Specific Interconnections

A. An Interconnection with Concurrency $\log_2 N$

To describe this interconnection, which was first presented in [10], let us initially consider $m = k$ 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".

1) Wiring: let $W(s, d)$ denote the transmitter ("0" or "1") used by s to reach d. (There is no choice of receiver in our case.) It is given by the following rule:

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

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 $k = 4$.

2) 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

d \ s	1 0 0 0	0 1 0 0	0 0 1 0	0 0 0 1
0 0 0 0	0	-0-	0	-0-
0 0 0 1	0	+0+	0	1
0 0 1 0	0	-0-	-1-	-0-
0 0 1 1	0	-0-	-1-	1
0 1 0 0	0	1	0	+0+
0 1 0 1	0	1	0	1
0 1 1 0	0	1	-1-	-0-
0 1 1 1	0	1	+1+	1
1 0 0 0	-1-	-0-	0	-0-
1 0 0 1	-1-	-0-	0	1
1 0 1 0	-1-	-0-	-1-	-0-
1 0 1 1	-1-	-0-	1	1
1 1 0 0	-1-	1	0	-0-
1 1 0 1	+1+	1	0	1
1 1 1 0	-1-	1	-1-	-0-
1 1 1 1	-1-	1	-1-	1

(a)

d \ s	1 0 0 0	0 1 0 0	0 0 1 0	0 0 0 1
0 0 0 0	1 0 0 0	0 1 0 0	0 0 1 0	0 0 0 1
0 0 0 1	1 0 0 1	(0 1 0 1)	0 0 1 1	0 0 0 0
0 0 1 0	1 0 1 0	0 1 1 0	0 0 0 0	0 0 1 1
0 0 1 1	1 0 1 1	0 1 1 1	0 0 0 1	0 0 1 0
0 1 0 0	1 1 0 0	0 0 0 0	0 1 1 0	(0 1 0 1)
0 1 0 1	1 1 0 1	0 0 0 1	0 1 1 1	0 1 0 0
0 1 1 0	1 1 1 0	0 0 1 0	0 1 0 0	0 1 1 1
0 1 1 1	1 1 1 1	0 0 1 1	(0 1 0 1)	0 1 1 0
1 0 0 0	0 0 0 0	1 1 0 0	1 0 1 0	1 0 0 1
1 0 0 1	0 0 0 1	1 1 0 1	1 0 1 1	1 0 0 0
1 0 1 0	0 0 1 0	1 1 1 0	1 0 0 0	1 0 1 1
1 0 1 1	0 0 1 1	1 1 1 1	1 0 0 1	1 0 1 0
1 1 0 0	0 1 0 0	1 0 0 0	1 1 1 0	1 1 0 1
1 1 0 1	(0 1 0 1)	1 0 0 1	1 1 1 1	1 1 0 0
1 1 1 0	0 1 1 0	1 0 1 0	1 1 0 0	1 1 1 1
1 1 1 1	0 1 1 1	1 0 1 1	1 1 0 1	1 1 1 0

(b)

Fig. 4. Wiring (a) and schedule (b) matrices (transposed) for an SDM-based $\mathcal{SHI}$ $k = 4$ SSs, each with two transmitters, and $2^k = 16$ DSs, each with a single receiver. In (a), only the transmitter number is shown. "+" and "-" highlight desired and stray transmissions in time-slot 5, respectively.

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 k-bit vectors, component by component; the result is treated as a k-bit number.) 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 [10] for generalizations and correctness proofs.

Fig. 4 illustrates the operation of the foregoing interconnection. It depicts its wiring and schedule matrices, highlighting the activity in time-slot number 5 (= 0101). (The matrices are transposed for formatting convenience.) In the wiring matrix (a), "+" is used to denote the (SS, DS) pairs that may communicate in this time slot, and a "$-$" in a row marks the corresponding DS as a stray destination, i.e., one that is not an addressee yet hears a transmission. Observe that whenever there is a "+" in a row, it is the only marked entry in that row. This means that a receiver that could be receiving a packet in this time slot indeed cannot hear any other transmissions, stray or otherwise; those would constitute a collision. Also, note that DS number 0101 doesn't hear anything in this time slot. This, in fact, is true of some station in every time slot and is the basis for a small improvement presented in [8] and explained briefly below.

Part (b) of the figure depicts the schedule matrix, with the "0101" entries highlighted. Clearly, these are the same entries that have the "+" in part 9a). Finally, the reader may relate the information in part (a) with the connections depicted graphically in Fig. 3.

The interconnection just described can be extended to the case of 2^k SS's ($m = n = N = 2^k$) 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 SS's 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 small improvement to this interconnection is presented in [8]. In every slot, it succeeds in utilizing the receiver that hears no transmissions. The interconnection described there employs $k+1$ types of SS's. It uses the same wiring function for the first k types; the last one chooses a transmitter based on the parity of the binary string constituting the id of the DS. A schedule is also presented there, but is omitted here for brevity.

III. Efficient Layout of the Interconnection with Concurrency $\log_2 N$

The only way to meet the implementation challenge is to overlap the splitting and merging of signal paths, so that each coupler in the path is balanced, i.e., 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 stretching the merging of signal paths with identical destinations across several stages of couplers.

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 $(N/k \times N/2)$ 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 $N/2 \cdot k = N \log_2 N/2$, 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. 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. Consequently, 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 stage-2 couplers. 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 $2 \cdot 2^{k/2}$; each has $2^{k/2}$ outputs, and a number of inputs that will be determined shortly.

Overlapping the splitting and merging. (Distributed merging.) We have thus far observed that 1) the signals of up to N/k transmitters may be merged, since they all have identical destinations, 2) 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 stage-1 couplers, we combine groups of x signals of 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 $z \cdot k/2$ 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 $z \cdot (k/2)$ 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 $x = 2^{k/2-1}$, which balances the stage-1 couplers. This yields $z = N/(kx) = 2^k/(k \cdot 2^{k/2-1}) = 1/k \cdot 2^{k/2+1}$. 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.

Final adjustments. 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. Example 2 below illustrates this situation.

A. Examples

Example 1: $k = 4; N = 2^4 = 16$. (See Fig. 5.)

There are 16 stage-3 couplers, one per DS, each of size (2×1).

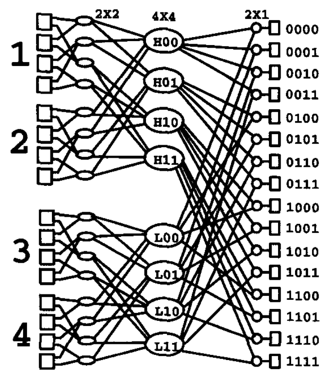

Fig. 5. Power-optimal layout for $k = 4$; $N = 2^4 = 16$. Rectangles denote stations, circles, and ellipses denote couplers. All couplers are balanced, except for the unavoidable (2 x 1) stage-3 couplers. The power split is 16, which is optimal.

There are $2 \cdot 2^{k/2} = 8$ stage-2 couplers. The first four represent all combinations of transmitter choices from the first $k/2 = 2$ "types" of SS's, and the remaining four represent the choices from the remaining two types of SS's.

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

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

Since we want $k/2 \cdot z = 2z = 4$, it follows that $z = 2$. Therefore, $x = \lceil N/(k \cdot z) \rceil = 2$. The resulting power split is $2 \cdot 4 \cdot 2 = 16$, which is optimal.

Fig. 5 depicts the resulting interconnection. Couplers are represented by circles or ellipses, and stations by rectangles. The SS's are **grouped by type**, with the number on the left denoting the type. Observe that each stage-1 coupler collects the signals of only one half of the stations of the same type, even though the signals of the same transmitters of all stations of the same type are going to the same receivers. The merger is completed in the stage-2 couplers. The benefit is that both stage-1 and stage-2 couplers are now balanced. Each stage-2 coupler is marked according to the source of the signals it carries. "H" stands for the $k/2$ "high-order" types, i.e., those that pick the transmitter to be used for any given connection based on one of the $k/2$ most significant bits in the destination's id. Similarly, "L" stands for "low-order" types. The two bits that follow the letter specify, for each of the $k/2$ types whose signals reach the coupler, which transmitter reaches the coupler. Finally, a stage-3 coupler is simply connected to the "H" and "L" couplers whose identifiers are equal to the $k/2$ most and least significant bits of the DS number, respectively.

The first example permitted a straightforward application of our "recipe". In the following example, we examine a situation in which this "recipe" yields non-integer numbers.

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

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 $2^{k/2} = 8$.

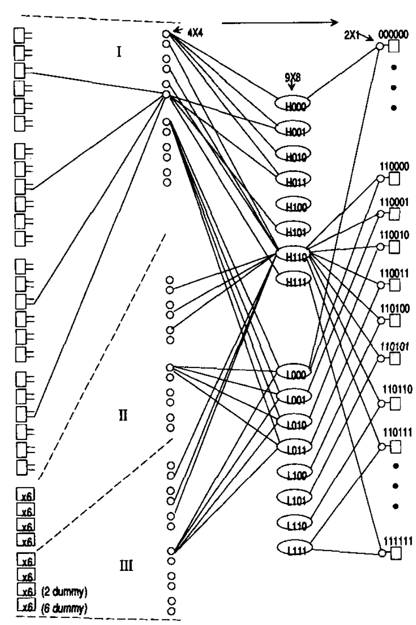

Fig. 6. Power-optimal layout for $k = 6$; $N = 2^6 = 64$. The power split is 72.

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

We want $k/2 \cdot z = 3z = 8$, but this produces a noninteger value for z. We therefore try the two nearest integer values, $z = 2$ and $z = 3$.

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:

- 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. 6, identified as "H110". 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 3: $k = 8; N = 2^k = 256$.

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.

IV. Efficient Layout of the Interconnection with Concurrency $\log_2 N + 1$

This interconnection, described in [8], appears to be more difficult to lay out efficiently because it is harder to discover symmetries when the wiring function for the SS's of type $(k + 1)$ differs 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 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.

V. Summary

The SDM-based SHI discussed in this paper permits $\log_2 N$ concurrent transmissions among N stations, each equipped with two transmitters and a single receiver or vice versa. Consequently, transmission rate need only be $1/\log_2 N$ of the aggregate network throughput. Alternatively, capacity is up to $\log_2 N$ times higher than that of a single channel for the same transmission rate and power budget.

We have shown how to implement this non-bus-oriented single-hop interconnection with 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 all but the (2×1) couplers in the last stage are balanced, it follows that the required number of elementary (2×2) couplers and fiber segments is approximately $N \cdot \log_2 N$, again optimal to within a constant. With unbounded coupler sizes, the number of fiber segments is $7N$, including the connections to transmitters and receivers, as compared with $3N$ for a maximum-concurrency bus-oriented SHI. Generalizations and extensions of the layout techniques presented here to other values of c_T and c_R will be presented elsewhere.

One apparent drawback of this interconnection is that round length is $N^2/\log_2 N$ and each pair of stations is allowed to communicate only during one slot per round. Therefore, although the capacity of the network is high, the maximum throughput between any two users is not. This suggests that with this operation scheme, the network is not suitable for applications in which a single (source, destination) pair needs a significant fraction of the network capacity. As was stated at the outset, however, the scenario that motivated this research is a large number of small stations that are being forced to transmit at a high rate merely due to a high aggregate throughput. Moreover, new schemes for operating this type of network are presently being explored, and are expected increase the flexibility of capacity allocation and reduce low-load delay. Finally, it should also be noted that the capacity that can be allocated to any single pair of stations when the traffic pattern is uniform drops as the number of stations increases. This, however, is inherent to any shared resource, and is in fact mitigated by the increase in concurrency.

In summary, then, by providing a power-efficient layout for an SDM-based interconnection with concurrency $\log_2 N$, we have shown that the shared directional multichannel can offer a significant advantage over buses, permitting the efficient construction of non-bus-oriented SHI's whose concurrency increases with network size while permitting the same transmission rate as a single bus interconnecting all stations, whose concurrency is only one. Moreover, the interconnections can be implemented using the simplest, non-tunable components. Finally, we note in passing that, as with a single shared channel, one can superimpose several networks of this type on the same physical medium using different wavelengths.

References

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

[2] 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, Mar. 1987.

[3] R. W. Boyd, *Radiometry and the Detection of Optical Radiation.* New York: John Wiley & Sons, 1983.

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

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

[6] C. A. Bracket, "Dense wavelength division multiplexing networks: Principles and applications," *IEEE J. Sleect. Areas Commun.,* vol. 8, no. 6, pp. 948–964, Aug. 1990.

[7] 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,* Mar. 1988.

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

[9] Y. Birk, "Fiber-optic bus-oriented single-hop interconnections among multi-transceiver stations," *J. Lightwave Technol.,* vol. 9, no. 12, pp. 1657–1664, Dec. 1991.

[10] 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, pp. 186–191, Jan. 1993.

[11] M. E. Marhic, "Hierarchic and Combinatorial Star Couplers," *Opt. Let.,* vol. 9, pp. 368–370, Aug. 1984.

Yitzhak Birk, photograph and biography not available at the time of publication.