

FIBER-OPTIC BUS-ORIENTED SINGLE-HOP INTERCONNECTIONS AMONG MULTI-TRANSCIVER STATIONS

Yitzhak Birk
Electrical Engineering Department
Technion – Israel Institute of Technology
Haifa 32000, Israel *birk@ee.technion.ac.il*

Abstract The single-path selective-broadcast interconnection (*SBI*) is a static, passive interconnection among a set of stations, each equipped with multiple, say c , transmitters and receivers. It employs c^2 buses, each interconnecting a subset of the stations, and provides a single common bus to any two stations. This paper explores those merits of this *SBI* which are related directly to its implementation in fiber optics. When compared with c buses, each interconnecting all stations, the *SBI* is shown to offer substantial advantages in power budget and the maximum number of stations that can be interconnected without repeaters or amplifiers. It is also attractive in terms of the required passive fiber-optic components such as fiber segments and star couplers. For a fixed power budget and direct detection, the capacity of this *SBI* is shown to be highest among bus-oriented single-hop interconnections for both a uniform traffic pattern and worst-case unknown skew.

Key words and phrases: *single-hop interconnections, fiber optic interconnections, bus-oriented interconnections, local area networks, FOLANs, selective-broadcast interconnections.*

1 Introduction

1.1 Background

A conventional local area network (LAN) uses a single shared channel to interconnect all stations. Notable examples are Ethernet [1] and radio networks. Consequently, the transmission rate must exceed both the required peak data rate for a single transmission and the aggregate throughput of all station pairs.

In the early days of LANs, the required network speed was indeed dictated by the peak single-user requirement. Lately, however, both the number of stations attached to a LAN and its usage by each station have been increasing rapidly. The increased usage is due to proliferation of distributed services, shared storage with diskless workstations, informa-

tion servers, distributed image-intensive applications, graphics terminals, etc., and is expected to grow even further. These changes are causing aggregate network throughput, not peak single-user requirements, to dictate the required transmission rate. Moreover, applications are unlikely to benefit from the higher transmission rate in low-load situations because of bottlenecks in their workstations. Users would thus be forced to pay for high-speed (expensive) hardware that is of little benefit to them, making shared channels less attractive.

Presently, transmission rate is decoupled from aggregate throughput by partitioning the network into multiple LANs, interconnected by routers and bridges. This solution is viable but expensive. Moreover, it places complex, active elements in the message path with negative implications on latency and reliability.

When LANs are implemented in fiber-optics, particularly with direct detection, power budget is an additional concern, and manifests itself as a limitation on the number of stations and/or the transmission rate. Alternatively, signals must be amplified and the network is no longer passive.

In view of the above, a question that arises naturally is what can be done to reduce the coupling between transmission rate and aggregate throughput and mitigate the power problem while retaining the simplicity and reliability of an "Ethernet-like" network, i.e., single-hop connectivity.

1.2 Single-hop interconnections

We define a *single-hop interconnection (SHI)* to be a static interconnection in which a message travels from the sender to the recipient without any intervention; i.e., no intermediate switches (as in multistage interconnections) and no need for forwarding by other stations (as in multi-hop networks). The interconnection network can thus be entirely passive. By "static" we mean that the set of receivers that can hear any given transmitter is fixed; frequency-agile transmitters and receivers, for example, are not permitted. Extreme instances of *SHIs* are a network with dedicated point-to-point links between every pair of stations, and a single broadcast

This work was done while the author was a Research Staff Member at IBM's Almaden Research Center.

bus (*SBB*). Notable uses of the *SBB* are Ethernet [1], radio networks and computer buses. *SHIs* are often desired due to their inherent reliability, low latency and simplicity in operation and maintenance.

1.3 The parallel broadcasts interconnection

With conventional signaling techniques, in which a bus can carry at most one successful transmission at any given time, and single-hop connectivity among all stations, each station must be equipped with multiple transmitters or receivers if any decoupling of transmission rate from aggregate throughput is to be achieved. The most obvious way of interconnecting user stations, each equipped with c transmitters and receivers, is to construct c subnetworks ("buses"), each interconnecting all stations through one of their transmitters and receivers [2],[3]. We refer to this as the *parallel broadcasts interconnection*, *PBI*. This would achieve a c -fold decoupling, but offers no advantage over *SBB* in hardware utilization and limited advantage in power budget.

1.4 The selective-broadcast interconnection

For a uniform traffic pattern (an equal amount of traffic between every pair of stations), one can do better than *PBI*. An important example is the single-path, unidirectional *selective-broadcast interconnection* (*SBI*) which comprises c^2 equally-populated buses such that any two stations have a common bus [4]. In [5], this interconnection was described and compared with *PBI* in terms of capacity and delay under an assumption of fixed transmission rate per bus, i.e., ignoring power budget. In [6], certain aspects of fiber-optic implementation were discussed. These included a method of using WDM to separate the various subnetworks, as well as an approach for combining WDM with space-division multiplexing. Also, a limited comparison of component requirements was carried out.

This paper focuses on the properties of *SHIs* in the context of fiber-optic implementations. It presents an extensive comparison between the single-path *SBI*, *PBI* and *SBB*. Fiber-optic component requirements and the maximum number of stations that can be accommodated without repeaters are compared under various assumptions and physical implementations. Also included is a capacity comparison assuming direct detection and fixed transmission power rather than fixed transmission rate. The single-path *SBI* is shown to dominate the others. (With fixed transmission rates, in contrast, the relative performance depends on the traffic pattern.)

The paper is organized as follows. In section II, we briefly

describe the single-path *SBI* as well as two methods of "interpolating" between an *SBI* and a *PBI*. This section is largely a summary of [5], and is included for completeness. Section III explores fiber-optic aspects of *SBI*, section IV discusses some of the results, and section V concludes the paper.

2 Construction of *SBI*s

2.1 Construction of a unidirectional, single-path *SBI*

Consider a set of N stations, each with c transmitters and c receivers. For simplicity of exposition, let us split each station into a *transmitting station* (TS) and a *receiving station* (RS). The TSs and RSs are partitioned into c groups of equal sizes. Next, c^2 subnetworks (buses) are constructed, such that subnetwork (i, j) connects the TSs of group i to the RSs of group j . More specifically, an RS uses its i th receiver to listen to TSs in the i th group of TSs; likewise, a TS uses its j th transmitter to reach RSs in the j th group of RSs. As depicted in Fig. 1, each TS has exactly one subnetwork

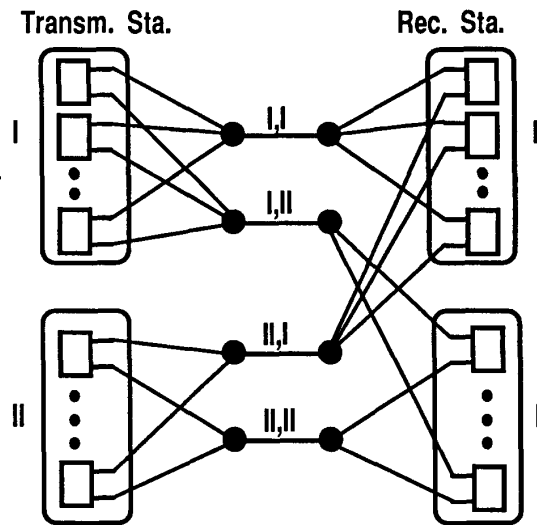


Figure 1: Single-path, unidirectional, equal-degree *SBI*. Several stations are shown for each group. $c = 2$; two groups; four subnetworks.

(bus) in common with any given RS. When $c = N - 1$, this *SBI* comprises a point-to-point link from each TS to each RS; when $c = 1$, it is an *SBB*. (*PBI*, in contrast, never becomes a collection of point-to-point connections.) Finally,

note that this can easily be generalized to the case of N_T transmitting stations and N_R receiving stations, each with c_T transmitters and c_R receivers, respectively. In this case, however, one cannot view a (TS, RS) pair as two parts of the same station.

2.2 Uniform-traffic capacity and station hardware

Let B denote the data rate of an individual transmission and thus the capacity of a single bus; C denotes the capacity of an entire interconnection.

For a Uniform traffic pattern and single-destination transmissions:

$$C^{\mathcal{P}BI} = c^{\mathcal{P}BI} \cdot B^{\mathcal{P}BI}; \quad (1)$$

$$C^{\mathcal{S}BI} = (c^{\mathcal{S}BI})^2 \cdot B^{\mathcal{S}BI}. \quad (2)$$

These expressions can be interpreted in several ways:

- With equal numbers of transmitters and receivers per station and equal transmission rates, the capacity of $\mathcal{S}BI$ is c times higher than that of $\mathcal{P}BI$.
- With equal capacities and the same numbers of transmitters and receivers per station, $\mathcal{S}BI$ can use slower (by a factor of c) and probably cheaper transmitters and receivers for the same aggregate throughput.
- With equal capacities and transmission rates, the required number of transmitters and receivers per station for $\mathcal{S}BI$ is only the square root of that for $\mathcal{P}BI$.

Since each subnetwork of $\mathcal{S}BI$ serves only N/c transmitting stations and N/c receiving stations, as compared with N in $\mathcal{P}BI$, the average utilization of transmitters and receivers can be c times higher than that of $\mathcal{P}BI$. In fact, this $\mathcal{S}BI$ is optimal in terms of uniform-traffic capacity, transmitter and receiver utilization, and power split (the number of receivers that hear a transmission, maximized over all transmitters) [4], [5], [7].

When all the traffic is from one TS group to one RS group, however, only one of the c^2 buses can be used, so the capacity of the single-path $\mathcal{S}BI$ drops to B , whereas that of $\mathcal{P}BI$ remains $c \cdot B$. Also, for any given source-destination pair, the maximum instantaneous data rate with $\mathcal{S}BI$ is B , as compared with $c \cdot B$ for $\mathcal{P}BI$. If the traffic pattern is known and static, one can assign stations to groups so as to balance the

load on the buses. Otherwise, our design goal is to maximize the uniform-traffic capacity C_{unif} subject to a guaranteed worst-case capacity C_{guar} . To achieve $C_{guar} = kB$, ($k \leq c$) an interconnection must provide k disjoint paths between any pair of stations.

2.3 k -path, unidirectional $\mathcal{S}BI$ s

In [3], it was shown that the maximum possible (relative) savings in transmitters and receivers without losing the flexibility of $\mathcal{P}BI$ is $(c+1)/N$, which becomes negligible as the number of stations increases. Given c , there thus appears to be a trade-off between C_{unif} and C_{guar} ; $\mathcal{P}BI$ and the single-path $\mathcal{S}BI$ are two extremes. Following are two parameterized compromises [4],[5].

Multiple single-path $\mathcal{S}BI$ s (MSP). k single-path $\mathcal{S}BI$ s are constructed, each of which utilizes $1/k$ of the transmitters and receivers of every station. Here,

$$C_{guar} = k \cdot B; \quad C_{unif} = \frac{c^2}{k} \cdot B; \quad \text{Power split} = k \frac{N}{c}. \quad (3)$$

A hybrid $\mathcal{S}BI$ - $\mathcal{P}BI$ interconnection. c' transmitters and c' receivers of each station are used for a $\mathcal{P}BI$, and the remaining ones are used for a single-path $\mathcal{S}BI$. See Fig. 2. Here,

$$C_{guar} = (c' + 1) \cdot B; \quad C_{unif} = (c' + (c - c')^2) \cdot B; \\ \text{Power split (worst case)} = N. \quad (4)$$

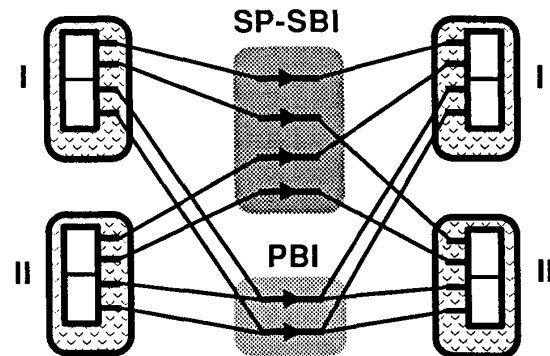


Figure 2: Unidirectional hybrid $\mathcal{S}BI$ - $\mathcal{P}BI$. $c=4$, $c'=2$. One station is shown for each group of $\frac{N}{2}$. $C_{guar} = 3$; $C_{unif} = 6$.

The hybrid outperforms the MSP , except for equality when $k \in \{1, c\}$, and offers more flexibility in allocation of hardware to the two components. On the negative side, the $\mathcal{P}BI$ part of the hybrid causes its utilization and power split to be no better than those of $\mathcal{P}BI$ [4].

2.4 Delay

For equal capacities, the delay with the single-path *SBI* was shown in [4],[5] to be higher than with *PBI*, which in turn was higher than that with a single bus. Also, the delay disadvantage of *SBI* in this case is smaller if the hardware savings (compared with *PBI*) take the form of fewer transmitters and receivers per station, rather than an equal number of slower ones. With equal transmission rates and equal numbers of transmitters and receivers per station, *SBI* still exhibits higher delay at low load, but much lower delay than the other two at higher loads; this is due to its higher capacity in this case.

2.5 Bidirectional SBIs

When stations are interconnected via a unidirectional *SBI*, the receivers of a station are connected to a different set of buses than its transmitters. Therefore, most existing access schemes cannot be used. In a bidirectional *SBI*, a station's transmitters and receivers must be connected to the same buses. A bidirectional *SBI* interconnecting a set of N stations, each with c transceivers, can be derived from a unidirectional *SBI* with $c + 1$ transmitters and receivers per station by merging the (i, j) bus with the (j, i) bus for all $i < j$, removing all (i, i) buses, and reducing the number of transmitters and receivers by one. The total number of buses is $c(c + 1)/2$. One can, however, obtain better results by applying the theory of block designs and projective geometry [8]. Specifically, whenever $(c - 1)$ is a power of a prime, it is possible to construct an *SBI* with

$$C_{bidirectional}^{SBI} = c^2 - c + 1. \quad (5)$$

For example, let $c=3$; divide the stations into 7 groups, numbered 0..6, and assign them to buses as follows: $\{0, 1, 3\}$, $\{1, 2, 4\}$, $\{2, 3, 5\}$, $\{3, 4, 6\}$, $\{4, 5, 0\}$, $\{5, 6, 1\}$, $\{6, 0, 2\}$. Fig. 3 illustrates the example. The design trade-offs and options are essentially identical to those of the unidirectional *SBI*. For further details, see [4],[5].

3 Fiber-Optic SBIs

In this section, we focus specifically on fiber-optic implementations of *SBIs*. The various issues are all related to power budget; reciprocity of star couplers is also taken into consideration. We begin by comparing this *SBI* with *PBI* in terms of the requirements for fiber-optic components and the maximum number of stations that can be accommodated with a given power budget and uniform-traffic capacity. We then




Figure 3: Bidirectional single-path *SBI*. $c = 3$; $c^2 - c + 1 = 7$ subnetworks. Each group is represented by a single station.

prove that with fixed transmission power on all buses, the single-path *SBI* is capacity-optimal and there is no trade-off between C_{guar} and C_{unif} . Unidirectional interconnections will be used for facility of exposition, but the results hold for bidirectional ones as well.

3.1 Passive fiber-optic component requirements

One might expect that the larger number of buses in *SBI* than in *PBI* would require more passive fiber-optic components, namely long fiber segments and directional star couplers. In this section, we show that this is not necessarily the case. We assume that fibers and couplers can operate at any transmission rate. The comparison is conducted for three sets of constraints: (i) equal B and C , (ii) equal c and C , and (iii) equal B and c . Two extreme configurations of an individual subnetwork are considered: a linear bus with taps and a centralized star.

Linear bus with taps. As shown in Fig. 4, each subnetwork is implemented as a single fiber that goes among the stations. Each transmitter is connected to this fiber by means




Figure 4: Linear-bus implementation of a subnetwork.

of a (2×2) star coupler, and the same is true of each receiver. The results are summarized in Table 1.

Centralized star. Here, a star coupler corresponds to a subnetwork, and a fiber corresponds to a transmitter or a receiver. The comparison of the interconnection component requirements is complicated by the fact that the required star

Conditions	Fibers		Couplers	
	<i>PBI</i>	<i>SBI</i>	<i>PBI</i>	<i>SBI</i>
equal:	<i>C/B</i>	<i>C/B</i>	$2N \cdot c_{PBI}$	$2N \sqrt{c_{PBI}}$
<i>C</i> and <i>B</i>	<i>C/B</i>	<i>C/B</i>	$2N \cdot c$	$2N \cdot c$
<i>c</i> and <i>B</i>	<i>c</i>	c^2	$2N \cdot c$	$2N \cdot c$
<i>C</i> and <i>c</i>	<i>c</i>	c^2	$2N \cdot c$	$2N \cdot c$

Table 1: Fiber-optic component requirements for a linear bus with taps.

couplers are of different sizes. We solve this by assuming that large couplers are implemented using small ones as building blocks [10]. (An $(M \times M)$ coupler can be constructed using $\frac{M}{p} \cdot \log_p M$ couplers of size $(p \times p)$, where p divides M .)

Table 2 summarizes the comparison. Perhaps the most interesting result is that for equal B and c , (the case in which *SBI* has higher capacity for identical active hardware,) and a star configuration, *SBI* requires fewer couplers and the same amount of fiber.

3.2 Maximum number of stations that can be accommodated (equal capacities and power)

Path loss is the ratio of the power at the output of a transmitter, P_T , and the power at the input of a receiver, P_R . Its constituents are:

- **Fan-out.** With direct detection and low-impedance optical detectors like those typically used for FOLANs at present, the reception of a signal “consumes” the power that is present at the receiver’s input, requiring a certain power level for reception. If a signal can reach several receivers, the level at each one is only a fraction of the transmitted power. This is in contrast with the case of coaxial cables and high impedance detectors, which sense the voltage and draw minimal amounts of power, or coherent optical detection.
- **Inefficient fan-in.** With an $(m \times n)$ lossless coupler made of fibers with constant cross section, the ratio of power at a single input to that at an output is $\max\{m, n\}$.¹

¹This is indirectly explained by the constant radiance theorem in optics [11], which states that when a narrow beam undergoes a linear lossless transformation, its radiance remains constant. A corollary of this is that the product of the cross-sectional area and the square of the numerical aperture of an optical beam must remain constant under any lossless linear transformation of that beam [12]. As a result, when several fibers are fused to form a single fiber, as is the case at the input of a star coupler, the cross-sectional area decreases and the numerical aperture increases. Unfortunately, the numerical aperture of the fiber

- **Excess loss.** This represents the imperfection of the coupler and its connectors.

We now redefine the *power split* of a path to be the path loss with ideal components (no excess loss). The power split of an interconnection is the maximum power split over all station pairs. This new definition does not alter the results of earlier sections, in which we took power split to be the number of receivers that can hear a transmitter.

Lemma 1. *The minimum power split for an *SHI* is N/c .*

Proof: c transmitters must jointly reach N stations. Thus, there must a transmitter that reaches at least N/c receivers. The single-path unidirectional *SBI* provides proof that this limit can be attained. \square

For a given transmitted power P_T , the maximum allowable path loss is $\frac{P_T}{P_{R_{min}}}$, where $P_{R_{min}}$ is the minimum amount of power required at the receiver. In studying the performance of existing optical receivers, it has been observed that over a wide range of transmission rates (100Mb/s–1Gb/s [13]; 155Mb/s–2.5Gb/s [14]), $P_{R_{min}}$ is roughly proportional to the transmission rate. This is consistent with a requirement of a minimal amount of energy per bit. Thus,

$$P_{R_{min}}(B) \approx B \cdot P_{R_{min}}(1). \quad (6)$$

The number of stations that can be accommodated by a passive fiber-optic interconnection is determined by the maximum path loss over all source-destination pairs. Since the subnetworks are disjoint, the first step in determining the maximum number of stations is to derive the maximum number per subnetwork (bus), N_b , as a function of the permissible path loss (*power margin*). Two configurations will be considered: a linear bus with taps, and a centralized star.

Maximum number of stations on a bus

- **Linear bus with taps.** A signal must first pass through a sequence of up to N_b couplers that collect the signals of downstream stations onto the bus, and then through one coupler for every receiver on the bus. Due to reciprocity of the couplers, the fraction of power that is coupled from a transmitter onto the bus is equal to the fraction that is taken off the bus to the dangling output of the coupler. This creates a trade-off in the selection of the coupling coefficient, and results in significant signal loss at each coupler [15]. This problem

is not any larger than that of the original one, so most of the power cannot propagate and is lost. The fact that the cross-sectional area again increases at the output of the coupler does not help.

Conditions	Fibers		Couplers	
	\mathcal{PBI}	\mathcal{SBI}	\mathcal{PBI}	\mathcal{SBI}
C and B	$2N \cdot c_{\mathcal{PBI}}$	$2N \sqrt{c_{\mathcal{PBI}}}$	$c_{\mathcal{PBI}} (N \times N)$	$c_{\mathcal{PBI}} \left(\frac{N}{c_{\mathcal{PBI}}} \times \frac{N}{c_{\mathcal{PBI}}} \right)$
c and B	$2N \cdot c$	$2N \cdot c$	$c (N \times N) = \frac{N \cdot c}{p} \log_p N (p \times p)$	$c^2 \left(\frac{N}{c} \times \frac{N}{c} \right) = \frac{N \cdot c}{p} \log_p \frac{N}{c} (p \times p)$
C and c	$2N \cdot c$	$2N \cdot c$	"	"

Table 2: Fiber optic component requirements, star implementation.

does not exist in the receiver couplers, each of which removes a small fraction of the signal from the bus. Nevertheless, the excess loss of a receiver coupler is compounded N_b times.

For simplicity of analysis, let us assume that all transmitter couplers have the same coupling ratio. Also, we take the effective transmitted power to be that which is actually coupled to the bus; finally, we lump the insertion loss of a receiver coupler together with the total loss of a transmitter coupler and the loss of signal that goes to the wrong output of each coupler and denote it $L (> 1)$. Thus, a signal travelling on the bus is attenuated by a factor L up to N_b times. The remaining loss is power split in receiver couplers and, if those are set to optimal ratios, is equal to N_b , the number of receivers on a bus.

The loss incurred by a signal from the first transmitter to the last receiver is thus

$$\frac{P_T}{P_R} \approx L^{N_b} \cdot N_b, \quad (7)$$

and the maximum number of stations on any given bus is such that

$$N_b + \log_L N_b \approx \log_L \left(\frac{P_T}{P_{R_{\min}}} \right). \quad (8)$$

This expression is clearly quite crude. Moreover, typical values of L , N_b and the power margin are such that the logarithmic term on the left hand side cannot be neglected. Nevertheless, (8) does offer some insight, telling us that for a linear bus with taps, the increase in N_b with an increase in power margin ($P_T/P_{R_{\min}}$) is sub-linear. Indeed, the change of N_b with power margin suggested here closely matches the numerical results in [15], which are based on a more detailed model.

- **Star configuration.** The star configuration is logically an $(N_b \times N_b)$ star. With the large star implemented using elementary $(p \times p)$ stars as building blocks, the signal passes through $\log_p N_b$ couplers on

its way from the transmitter to any receiver. The path loss is therefore

$$\frac{P_T}{P_R} = N_b \cdot L^{\log_p N_b} = N_b^{(1+\log_p L)}, \quad (9)$$

and the maximum number of stations on a bus is

$$N_b = \left(\frac{P_T}{P_{R_{\min}}} \right)^{\frac{1}{1+\log_p L}} \approx \frac{P_T}{P_{R_{\min}}} \approx \frac{P_T}{P_{R_{\min}}(1) \cdot B}. \quad (10)$$

Comparison of N_b among \mathcal{SBI} , \mathcal{PBI} and \mathcal{SBB}

We assume equal capacity C for all three, and equal c for \mathcal{SBI} and \mathcal{PBI} . As a result, \mathcal{SBI} can use a lower transmission rate.

- **Linear bus with taps.** Since we do not have a precise quantitative formula, let us consider the specific example of a LAN with an aggregate capacity $C = 900\text{Mb/s}$; $P_T = 1\text{mW}$; minimum energy per bit (at the receiver) is $1.5 \cdot 10^{-15}$ Joules (20dB above the quantum limit); $c = 3$. Results for coupler losses of 0.5dB and 1.0dB are presented in Table 3, which also depicts the maximum total number of stations. (Coupler loss includes connections, excess loss and fiber loss.) The results for 1dB were taken from Fig. 6 in [15]; those for 0.5dB were obtained using (8), with L chosen to match the result in Fig. 7 [15] for a 40dB power margin (PM).
- **Star configuration.** Let N_0 denote the maximum number of stations that can be accommodated by the \mathcal{SBB} with capacity C . It follows from (10) that

$$N_b^{\mathcal{PBI}} \approx c \cdot N_0 \quad (11)$$

$$N_b^{\mathcal{SBI}} \approx c^2 \cdot N_0 \quad (12)$$

Maximum total numbers of stations

- **Linear bus with taps:** Numerical results are presented in Table 3. For $c = 3$, \mathcal{SBI} offers an advantage

Topology	B (Mb/s)	PM (dB)	N_b		N (total)	
			0.5dB	1.0dB		
<i>SBB</i>	900	28.7	17*	11	17*	11
<i>PBI</i>	300	33.4	20*	13	20*	13
<i>SBI</i>	100	38.0	24*	15	72*	45
		40.0	27	16		

Table 3: Maximum number of stations — linear bus with taps. Values of N_b and N marked with “*” are based on (8); the others are from [15], Fig. 6,7. ($c = 3$.)

by more than a factor of three. Moreover, since the benefit is due primarily to the fact that $N = c \cdot N_b$ (for *SBI*), the results would remain similar if we used equal transmission rates.

• **Star configuration:**

$$N^{\mathcal{PBI}} = c \cdot N_0. \quad (13)$$

$$N^{\mathcal{SBI}} = c^3 \cdot N_0; \quad (14)$$

The maximum number of stations which can be accommodated by *SBI* is thus always higher than the corresponding numbers for the single bus or *PBI* by at least a factor of c , due to the fact that $N^{\mathcal{SBI}} = c \cdot N_b^{\mathcal{SBI}}$. An additional advantage of up to c^2 over the single bus and up to c over *PBI* is a by-product of the reduced transmission rate.

3.3 Capacity with equal power

Our discussion of ways of accommodating variable or unknown traffic patterns in section II was based on an assumption of fixed transmission rate on a bus, and identified a trade-off between guaranteed worst-case capacity and uniform-traffic capacity. We now revisit the proposed compromises under an assumption of equal transmission power, which is the common case with fiber-optic implementations. (Transmission rate will depend on power budget.) The analysis will first be carried out for a star implementation of each subnet, and then for a linear bus with taps.

Star. With each subnet implemented as a star, the power at each receiver is inversely proportional to the number of receivers on the bus. (This remains true for lossy components, since the number of those in any given path is logarithmic in the number of receivers.) For fixed transmission power, the maximum transmission rate is therefore inversely proportional to the number of receivers on the bus.

Let B_0 denote the maximum transmission rate on a bus with N stations. We now recompute the capacities of the different configurations.

k -path *SBI* with equally populated buses. (E.g., *MSP*.) The number of receivers on each bus is $\frac{N}{c/k}$, so the permissible transmission rate is $\frac{c}{k} B_0$. Thus,

$$C_{guar} = k \cdot 1 \cdot \frac{c}{k} B_0 = c \cdot B_0; \quad (15)$$

$$C_{unif} = k \cdot \left(\frac{c}{k}\right)^2 \cdot \frac{c}{k} B_0 = \frac{c^3}{k^2} \cdot B_0. \quad (16)$$

Hybrid *SBI-PBI*. The permissible data rate on each bus of the *SBI* portion is $(c - c') \cdot B_0$, but on the *PBI* buses it is only B_0 . Therefore,

$$C_{guar} = 1 \cdot (c - c') \cdot B_0 + c' \cdot 1 \cdot B_0 = c \cdot B_0; \quad (17)$$

$$C_{unif} = ((c - c')^3 + c') \cdot B_0. \quad (18)$$

Surprisingly, C_{guar} is identical in all cases, so we are free to optimize for C_{unif} . This is attained with $k = 1$ or $c' = 0$, both of which correspond to the single-path *SBI*. Moreover, while we implicitly permitted different transmission rates on different buses, the optimal topology does not exploit this! We conclude that the inclusion of the interplay between the number of stations on a bus and the allowable transmission rate strongly favors the *SBI*. For example, the uniform-traffic capacity of a single-path *SBI* with $c = 2$ would be (at least) four times higher than that of a *PBI* with $c = 2$ and the same power budget. (The worst-case capacities are equal.)

Linear bus with taps. In determining the maximum number of stations on a bus, we noted that a three-fold increase in power margin did not substantially increase N_b . In general, a large difference in power margin would be required to change N_b by even a small integral factor. In the present discussion, N and c are equal for all interconnections, so the number of stations on each of the *SBI* buses is smaller than those of *SBB* and *PBI* by a factor of c . Consequently, we expect a very large change in the power margin, which in

turn would result in a similarly large change in maximum transmission rate (based on (6)) and thus in capacity.

As an example, we again use numbers from Fig. 6 in [15]. (Unfortunately, we cannot use the same example as before because the numbers fall off the curves.) Reducing N_b from 20 to 10 (corresponding to $c = 2$) changes the required power margin from 48dB to 28dB. Under these conditions, the permissible transmission rate with the single-path *SBI* would therefore be 100 times higher than with *PBI* or *SBB*. Thus, even if only one of its four buses could be used due to traffic skew, the *SBI*'s capacity would still be 50 times higher than that of *PBI*. With $c = 3$ the results would be even more dramatic.

Power budget (and thus transmission rate) is least sensitive to the number of stations on a bus in the star implementation. Since the optimal solution for the star was the one with the fewest stations per bus, it is therefore clearly optimal for any other implementation. We summarize this in the following theorem:

Theorem 2. *Given N stations, each with c transmitters and receivers, fixed transmission power and required energy per bit, and a required guaranteed capacity (over the entire range of traffic skews), the single-path *SBI* has the highest uniform-traffic capacity of all static, passive, single-hop fiber-optic interconnections. Moreover, the capacity of this *SBI* is greater than or equal to that of any other *SBI*, *PBI* or combination thereof for any traffic pattern. □*

Note. The careful wording of the theorem reflects the fact that for a known sparse traffic pattern, one can sometimes construct a single broadcast bus (to guarantee single-hop connectivity) along with a collection of point-to-point links between pairs of stations that communicate extensively, thereby achieving a very high capacity for that specific pattern. The reader should also note that while the *SBI* has the same guaranteed capacity as *PBI*, the degree to which transmission rate is decoupled from capacity with *SBI* still depends on the traffic pattern. Specifically, when all the traffic uses the same bus, there is no decoupling.

4 Discussion

Having established various advantages of *SBI*, in this section we revisit some of the costs and apparent disadvantages. Also, a number of issues that are outside the main thrust of this paper but may be of interest to the reader are discussed briefly.

The cost of multiple transmitters and receivers. We have shown that the single-path *SBI* offers substantial advantages over a single broadcast bus or even multiple broadcast buses. However, one may still wonder about the cost of multiple transmitters and receivers per station. Although each station requires several network adapters, these adapters can be much slower (for equal capacities) and cheaper. In fact, there is always a speed beyond which several slow adapters would cost less than a single fast one. The break-even point for *SBI* is lower than that for *PBI* due to the sharp increase in capacity with an increase in the number of adapters. If one were willing to design special multi-adapters, further substantial savings would be attained.

Peak instantaneous rate. A perceived disadvantage of *SBI* relative to *PBI* is that *PBI* can make its entire capacity available to a single pair of stations whereas *SBI* cannot. However, the reduced number of stations on an *SBI* bus permits a transmission rate that is at least c times higher than on a single *PBI* bus. Stated differently, the capacity of a single *SBI* bus under an equal-power constraint is equal to or even greater than that of the entire *PBI*. The same is true for a comparison of *SBI* with *SBB* with equal transmission power per station. It is also worth noting that using the entire capacity of *PBI* for a single message complicates the protocols and requires packet reassembly at the destination.

Fault tolerance. The single-path *SBI*, unlike *PBI*, provides only a single path between any pair of stations. This path constitutes a single point of failure. However, multi-hop communication could be used in case of failure. With 2-hop routing, the interconnection can tolerate any $c - 1$ faults, like *PBI*.

Alternative implementations.

- **Rings.** High-speed LANs are often implemented as rings rather than buses. While power-budget advantages are no longer relevant, *SBI* retains some of its other advantages. A similar observation applies to systems with other forms of signal amplification or coherent detection.
- **Spatial/spectral subnetwork separation.** Figures 1 and 3 imply a spatial separation between the subnetworks, and call for c physical transmitters and receivers per station. Nevertheless, separation can also be achieved in the frequency domain, polarization, angle [16] (when relevant) and others, and the actual number of transmitters per station can sometimes be as low as one. It is also possible to combine different separation methods. For example, one could combine spatial and spectral separation so that any two subnetworks are separated in space, wavelength or both. The interested reader is referred to [17],[6] and [4] for a de-

tailed discussion if this issue, including an algorithm for assigning wavelengths to subnetworks and the possible savings in fiber-optic components.

- **Spread-spectrum.** Decoupling of transmission rate from aggregate throughput was one of the goals of *SBL*. Further decoupling can be attained through the use of code-division multiple access [18][19][20] in the implementation of the individual buses. CDMA should thus be viewed as complementing *SBL* rather than competing with it.
- **The impact of using real channel access schemes.** Most channel access schemes operate more efficiently at lower transmission rates [9]. Consequently, the fact that the total network capacity is divided among more buses permits more efficient operation. The use of real access schemes thus has a favorable effect on *SBL*'s merits relative to those of *SBB* or *PBL*.

5 Conclusions

Equipping every station on a LAN with a small number of transmitters and receivers and interconnecting the stations through a collection of buses such that any two stations have a single bus in common can result in a sharp increase in total network capacity. With fiber-optic implementations, additional benefits include a larger number of stations for given capacity and power, as well as other important benefits. For a fixed transmission rate, there is a trade-off between uniform-traffic capacity and guaranteed capacity (over the range of traffic pattern). For a fixed power budget, the capacity of the single-path *SBL* is at least as high as those of other topologies in the worst case, and is much higher for a uniform traffic pattern. *SBLs* can be operated using existing network adapters and protocols.

SHIs cannot compete with multi-stage interconnections or with multi-hop ones in terms of performance; nevertheless, this paper helps demonstrate that their performance can be extended quite dramatically beyond that of a single bus while retaining the simplicity and reliability of single-hop communication through a purely passive communications fabric. It is also worth noting that a much higher capacity can be attained by using frequency-agile transmitters; the interconnection, however, would no longer be static.

Our discussion was restricted to bus-oriented *SHIs*. While the capacity of these increases with c , it does not grow with an increase in the number of stations. With unidirectional media, such as fiber optics with directional star couplers, more general *SHIs* can be constructed [4], whose capacity can also grow with N [21].

Acknowledgment. I wish thank Michel Marhic and Mehdi Nassehi for reviewing portions of this paper.

References

- [1] R.M. Metcalfe and D.R. Boggs, "Ethernet: distributed packet switching for local computer networks", *Communications of the ACM*, vol. 19, no. 7, pp. 395-403, 1976.
- [2] M.A. Marsan and D. Roffinella, "Multichannel local area network protocols", *IEEE JSAC*, vol. SAC-1, no. 5, pp. 885-897, Nov. 1983.
- [3] T. Lang, M. Valero and M.A. Fiol, "Reduction of connections for Multibus Organization," *IEEE Trans. Comp.*, vol. C-32, no. 8, pp. 707-715, Aug. 1983.
- [4] 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.
- [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] Y. Birk, F.A. Tobagi and M.E. Marhic, "Selective-broadcast interconnections (*SBL*) for wideband fiber-optic local area networks", *Proc. SPIE conference on Fiber Optic Broadband Networks*, Cannes, France, Nov. 1985.
- [7] M.T. Busche, "Throughput, delay and power losses in bus-oriented optical networks with multiple transmitters and receivers per station", University of Illinois at Urbana-Champaign Report UILU-ENG-91-2205, Jan. 1991.
- [8] M. Hall, Jr., *Combinatorial Theory*, Blaisdell Publishing company, 1967.
- [9] F.A. Tobagi, "Multiaccess protocols in packet communication systems," *IEEE Trans. Commun.*, vol. COM-28, Apr. 1980.
- [10] M.E. Marhic, *Opt. Lett.*, vol. 9, p. 368. 1984.
- [11] R.W. Boyd, *Radiometry and the Detection of Optical Radiation*, John Wiley & sons, New York, 1983.
- [12] J.W. Goodman, "Fan-in and fan-out with optical interconnections," *Optica Acta*, vol. 32, no. 12, pp. 1489-1496, 1985.

- [13] T.V. Muoi, "Receiver design for high-speed optical-fiber systems," *IEEE J. Lightwave Technology*, vol. 2, pp. 243-267, June 1984.
- [14] C.A. Bracket, "Dense Wavelength Division Multiplexing Networks: Principles and Applications", *IEEE JSAC*, vol. 8, no. 6, pp. 948-964, Aug. 1990.
- [15] M.M. Nassehi, F.A. Tobagi and M.E. Marhic, "Fiber optic configurations for local area networks," *IEEE JSAC*, vol. SAC-3, no. 6, pp. 941-949, Nov. 1985.
- [16] R.C. Stearns, C.K. Asawa and S.K. Yao, "Angular division multiplexer for fiber communication using graded-index rod lenses," *IEEE J. of Lightwave Technology* vol. 2, p. 358. 1984.
- [17] M.E. Marhic, Y. Birk and F.A. Tobagi, "Selective broadcast interconnection (*SBT*): a novel scheme for fiber-optic local area networks", *Optics Letters*, Dec. 1985.
- [18] Special issue on spread-spectrum communication, *IEEE Trans. Commun.*, vol. COM-30, no. 5, pp. 817-1072, May. 1982.
- [19] J.Y. Hui, "Pattern Code Modulation and Optical Decoding - A Novel Dode-Division Multiplexing Technique for Multifiber Networks," *IEEE J. Sel. Areas in Commun.*, SAC-3, No. 6, Nov. 1985.
- [20] Y.L. Chang and M.E. Marhic, "2nd Codes for Optical CDMA and Associated Networks", *Conf. Dig. IEEE LEOS/COMSOC Summer Topical on Optical Multiple Access Networks*, Monterey, CA, Jul. 1990.
- [21] 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.