Selective Broadcast Interconnection (sbi) For Wideband Fiber-Optic Local Area Networks

Birk, Yitzhak, Marhic, Michel, Tobagi, Fouad

Yitzhak Birk, Michel E. Marhic, Fouad A. Tobagi, "Selective Broadcast Interconnection (sbi) For Wideband Fiber-Optic Local Area Networks," Proc. SPIE 0585, Fiber Optic Broadband Networks, (6 July 1986); doi: 10.1117/12.966543

Event: 1985 International Technical Symposium/Europe, 1985, Cannes, France

Selective broadcast interconnection ($\mathcal{SBI}$) for wideband fiber-optic local area networks

Yitzhak Birk
Michel E. Marhic*
Fouad A. Tobagi

Computer Systems Laboratory, Departments of Electrical Engineering and Computer Science
Stanford University, Stanford, CA 94305, U.S.A.

Abstract

The selective broadcast interconnection ($\mathcal{SBI}$) is a scheme for directly interconnecting transmitting stations, each equipped with C_T transmitters, and receiving stations, each equipped with C_R receivers, such that each transmitting station is always connected to all receiving stations through passive communication channels with no intermediate switches. $\mathcal{SBI}$ consists of $C_T \cdot C_R$ separate broadcast subnetworks, each of which interconnects a subset of transmitting stations and a subset of receiving stations, such that each transmitting station and receiving station are interconnected through a single subnetwork. Each subnetwork is shared by its transmitting members via some multiple-access scheme.

Comparing $\mathcal{SBI}$ with $C_T{=}C_R{=}C$ with the use of C broadcast buses, each connecting all transmitting stations to all receiving stations, one finds that in some cases, including that of equal single-destination traffic requirements for all source-destination pairs, the aggregate throughput with $\mathcal{SBI}$ can be higher by a factor of C, while the stations' hardware is the same. For nonuniform traffic requirements, however, the maximum aggregate throughput with $\mathcal{SBI}$ can be C times lower (in extreme situations). For fiber-optic implementations employing a central wiring closet, the two schemes require the same amount of fiber and, if the same elementary couplers are used to construct the required star couplers, $\mathcal{SBI}$ requires fewer couplers. Clearly, the same number of couplers and up to C times more fibers may be required for $\mathcal{SBI}$ in a linear-bus implementation. In all cases, transmitter power need only reach N/C receivers with $\mathcal{SBI}$ (instead of N with C parallel buses); this allows to accommodate a larger number of stations when implementing the interconnection with passive components.

Introduction

Many configurations considered to date for fiber-optic local area networks (FOLAN's) are of the conventional broadcast type, in which N stations, each with a single transmitter and a single receiver of peak data rate B, use time-division multiplexing (TDM) to share a common communication channel (bus).[1,2] Various access schemes can be used to control the use of this channel.[3,5] In such networks, there can be at most one ongoing successful transmission; consequently: (i) the aggregate throughput S cannot exceed B, and (ii) each station must be capable of transmitting and receiving at a peak rate in excess of S, regardless of its own throughput requirements.

Broadcast FOLAN's are also subjected to a power limitation when implemented with passive interconnection components: the output of each transmitter, of power P_T, must reach each receiver, of sensitivity P_R, with a level of at least P_R. For conventional broadcast FOLAN's, this translates into the requirement that $N < P_T/P_R$. Since an increase in B is typically accompanied by a linear increase in P_R,[6] the maximum N that can be accommodated by such FOLAN's for a given P_T decreases with an increase in B, and can become quite small at high speeds. The power problem is most severe if a simple but energy-inefficient structure, such as a linear bus, is used.[7] Power limitations can be alleviated by the use of repeaters, but this adds to the cost and often decreases reliability.

Most other FOLAN's considered to date are of the ring or loop type, which employ point-to-point optical links and are therefore not subject to significant power limitations[8]; nevertheless, the restrictions on B and S are similar to those of broadcast networks: although several transmissions can coexist on different sections of these systems, S is still limited to being less than $2B$ (under uniform traffic requirements).

---

*On leave from the Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, Illinois 60201.

In order to increase S beyond B or $2B$, one must use systems with some degree of concurrency; i.e., systems that can accommodate several ongoing successful transmissions. To achieve this with point-to-point optical links, one could use switching nodes (some or all of which might be the stations themselves), each equipped with several transmitters and receivers, that route messages according to their destinations. Several such methods have been proposed for FOLAN's,[9,10] and many more could be adapted to them when faster optical and/or electronic switching and electronic processing become commonplace.

In this paper, we limit the discussion to networks in which all stations are directly interconnected through passive communication channels, without any intermediate switches; i.e, there is always a communication path between each pair of stations. Furthermore, we focus on the use of such networks for single-hop communication. At the outset, we observe that the only way of achieving concurrency in such networks is through the use of multiple transmitters and/or multiple receivers at each station.

A straightforward approach to achieving concurrency is to employ C broadcast subnetworks, each connecting all stations using one of their C transmitters and receivers; these subnetworks could either be used independently for bit-serial transmission, or in parallel for the transmission of C-bit words. In either case, each station must have C transmitters and C receivers, and $S \leq C \cdot B$. We denote this scheme by $\mathcal{PBI}$, for parallel broadcast interconnection. Since $\mathcal{PBI}$ is essentially a replication of a conventional broadcast network, its throughput is C times that of such a network; however, power budget is the same. Note that $\mathcal{PBI}$ provides C independent paths between each pair of stations.

The selective broadcast interconnection ($\mathcal{SBI}$)

We now propose a scheme in which, unlike $\mathcal{PBI}$, each subnetwork interconnects transmitters belonging to a subset of the stations with receivers belonging to only a (possibly different) subset of the stations, in a way that provides a single transmission path from each station to each other station; whenever a station wishes to transmit to some other station, it does so using the transmitter that is heard by a receiver of the intended recipient. In other words, a station selectively broadcasts to a destination-dependent subset of stations; hence the name $\mathcal{SBI}$. Since each subnetwork in $\mathcal{SBI}$ interconnects fewer stations than in $\mathcal{PBI}$, it follows that for a given number of transmitters and receivers per station, there can be more subnetworks and hence a potentially higher level of concurrency. In fact, if all subnetworks are of the same size, $\mathcal{SBI}$ has the maximum possible number of subnetworks.

In order to illustrate the versatility of $\mathcal{SBI}$, we consider the somewhat more general problem of connecting N_T transmitting stations, each with C_T transmitters, to N_R receiving stations, each with C_R receivers. While $N_T \neq N_R$ is also possible with $\mathcal{PBI}$, $C_T \neq C_R$ is not; therefore, whenever comparing the two, we will assume that $C_T = C_R = C$. For convenience , we assume that

$$N_T = P \cdot C_R \qquad \text{and} \qquad N_R = Q \cdot C_T \tag{1}$$

where P and Q are integers. The common case wherein $N_T = N_R$, $C_T = C_R$, and transmitting and receiving stations are paired to form bidirectional stations, is a special case of the above.

To construct $\mathcal{SBI}$, we arrange the transmitting stations in C_R disjoint groups, each with P stations; similarly, we arrange the receiving stations in C_T disjoint groups, each with Q stations. Next, we construct $C_T \cdot C_R$ subnetworks, each connecting a group of transmitting stations to a group of receiving stations. Viewed differently, each transmitting station uses its jth transmitter to send messages to the jth group of receiving stations; similarly, each receiving station uses its ith receiver to receive messages from the ith group of transmitting stations. Fig. 1, which shows the transmitting and receiving stations at opposite ends of the drawing, represents a logic diagram of the connections; observe that each transmitting station has only one subnetwork in common with any given receiving station. Since there are $C_T \cdot C_R$ disjoint subnetworks, the degree of concurrency in this arrangement can reach the value $C_T \cdot C_R$. Finally, we note that $\mathcal{SBI}$ reduces to well-known configurations in the following limiting cases:
a) $C_T = C_R = N_T = N_R = N$; this corresponds to a fully connected topology with a point-to-point link from each transmitting station to each receiving station.
b) $C_T = C_R = 1$; this is a conventional broadcast network (single channel).

Performance

While the concurrency provided by SBI is $C_T \cdot C_R$, the extent to which this concurrency can be utilized depends on the traffic requirements. Therefore, we will always evaluate the performance of SBI in conjunction with those requirements, and whenever we mention concurrency it should be interpreted as the "effective" concurrency for those requirements. (The effective degree of concurrency provided by PBI is always C.)

Throughput and the required numbers of transmitters and receivers

Uniform single-destination traffic requirements. Assuming single-destination packets and equal traffic requirements for all source-destination pairs, (also referred to as "uniform traffic requirements",) the $C_T \cdot C_R$ subnetworks of SBI can be treated as independent, identical conventional broadcast networks, each connecting P transmitting stations to Q receiving stations; the throughput of SBI can then be summarized by the expression

$$S^{SBI} = K^{SBI} \cdot C_T^{SBI} \cdot C_R^{SBI} \cdot B^{SBI} \tag{2}$$

Since PBI can always be treated as C^{PBI} independent, identical conventional broadcast networks, each connecting N_T transmitting stations to N_R receiving stations, its performance can be summarized by the expression

$$S^{PBI} = K^{PBI} \cdot C^{PBI} \cdot B^{PBI} \tag{3}$$

K is a constant which depends on the channel access scheme ($0 < K \leq 1$). Also, for a given access scheme, it may vary with B and with the number of transmitters per subnetwork: for access schemes such as CSMA[3] and Expressnet[4], that require a time gap between consecutive packets and groups of packets, respectively, which is on the order of the end-to-end propagation delay through the network, the fraction of channel time that is taken up by the gaps increases as B increases (and hence packet transmission time decreases), thus decreasing K. This problem is less severe in other schemes, such as TDMA, although one could argue that the gap is contained in the slot size. Finally, for pure ALOHA[3], there is no dependence on B. For some random-access schemes, such as ALOHA, K increases as the number of contending transmitters decreases; for others, such as CSMA, there is no significant dependence, and for Expressnet as well as several other round-robin schemes,[5] K actually increases as the number of active transmitters increases.

Assuming that $K^{PBI} = K^{SBI}$, and, to permit comparison, that $C_T{=}C_R$ and that N_T and N_R are the same in both systems, these expressions can be interpreted in several different ways:

- With $C^{SBI} = C^{PBI} = C$ and $B^{SBI} = B^{PBI}$, the aggregate throughput for SBI is C times higher than for PBI (since it increases quadratically rather than linearly with C).
- With $S^{SBI} = S^{PBI}$ and $C^{SBI} = C^{PBI} = C$, the transmission rate required with SBI is C times lower than that required with PBI; i.e., slower (and cheaper) transmitters and receivers may be used for the same throughput.
- With $S^{SBI} = S^{PBI}$ and $B^{SBI} = B^{PBI}$, $C^{SBI} = \sqrt{C^{PBI}}$; i.e., SBI requires fewer transmitters and receivers.

Finally, we note that since each subnetwork of SBI serves only N/C transmitting stations and N/C receiving stations, as compared with N in PBI, the average fraction of time that a subnetwork of SBI serves each of its member stations is higher by a factor of C than that fraction with PBI. It follows that the average utilization of transmitters, of receivers and, in the case of fiber-optic implementations with a central wiring closet, of the fibers connecting stations with the wiring closet, is also higher by the same factor.

Other traffic requirements . For nonuniform single-destination traffic requirements, the throughput with SBI may become as low as that with a single broadcast network. This happens, for example, if all the traffic is from a single group of transmitting stations to a single group of receiving stations, in which case only one subnetwork can be used. We also note that, for any given source-destination pair, the maximum instantaneous data rate with SBI is B, as compared with $C \cdot B$ with PBI. For multi-destination packets, we observe that multicast to any subset of receiving stations that are connected to a given subnetwork is a byproduct of any successful transmission over that subnetwork. However, when several of a node's transmitters must transmit in order to cover the set of intended recipients, the performance of SBI degrades, and if transmissions by all C_T transmitters are required, as is the case for full broadcast, SBI loses its throughput advantage over PBI.

Power

An important aspect in which SBI always outperforms PBI is power splitting. While the use of PBI requires splitting the power of each transmitter N_R ways, it suffices to split it $N_R/C_T{=}Q$ ways for SBI. Observe that if a station has C_T transmitters and is to be connected to N_R different receiving stations, the power of each transmitter must be split at least N_R/C_T ways. Consequently, SBI is optimal in this sense and no other broadcast interconnection can do better.

Implementing SBI

In this section we discuss issues pertaining to fiber-optic implementation of SBI. Whenever applicable, a comparison will be made with PBI. Whenever we use C, it is implied that $C_T{=}C_R{=}C$.

Centralized fiber-optic implementations of the interconnection

Space-division multiplexing (SDM). Fiber and coupler requirements:

- PBI and SBI require the same number of fibers for connecting the stations with the central location (assuming the same C).
- As seen in Fig. 1, SBI requires $C_T C_R$ $(P \times Q)$ couplers, one for each subnetwork. When $C_T{=}C_R{=}C$, this translates to C^2 $(\frac{N}{C} \times \frac{N}{C})$ couplers, whereas PBI requires C $(N \times N)$ couplers. While a cost comparison is generally quite complicated, we note that if $N_T{=}N_R{=}p^n$, $C_T{=}C_R{=}p^m$, $(p, m, n$ integers) and one uses $(p \times p)$ elementary couplers as building blocks,[11] then fewer couplers are required for SBI:

$$\text{Number of } (p \times p) \text{ couplers required for } SBI = \frac{NC}{p} \cdot \log_p \left(\frac{N}{C}\right) \tag{4}$$

$$\text{Number of } (p \times p) \text{ couplers required for } PBI = \frac{NC}{p} \cdot \log_p N \tag{5}$$

We conclude that, for centralized implementations, the cost of interconnection components is not higher for SBI than for PBI.

Wavelength-division multiplexing (WDM). Wavelength division multiplexing can replace SDM in the implementation of SBI (as well as PBI). Rather than discuss it as a separate issue, we will next discuss hybrid SDM-WDM implementations, of which WDM is a special case.

Hybrid SDM–WDM implementation. We now show how WDM with W wavelengths can be used in conjunction with SDM to reduce the number of fibers and couplers required for the implementation of SBI. (This can also be done for PBI, in which case the use of WDM is straightforward.) We start out by outlining the scheme for reducing the number of fibers, then proceed to explore the wavelength assignment problem and, lastly, discuss the possible reduction in the number of couplers.

To reduce the number of fibers connected to each station, while, for the time being, using a separate star coupler for each subnetwork, we proceed as follows. At each transmitting station, outputs of transmitters that use different wavelengths are multiplexed onto a common fiber, thus reducing the number of fibers leading from each transmitting station to the wiring closet. At the wiring closet, demultiplexing takes place, yielding the $N_T \cdot C_T$ distinct signals from all transmitting stations. These are then fed into individual inputs of the couplers, as before. At the output of the couplers, signals of different wavelengths which are destined to the same receiving station are multiplexed onto a common fiber, thereby reducing the number of fibers between the wiring closet and each receiving station. At each receiving station, demultiplexing again takes place, and the single-wavelength signals are connected to individual receivers.

Since each of the $C_T \cdot C_R$ couplers represents a subnetwork, it follows that each coupler must carry a single wavelength. Consequently, we will assign wavelengths to couplers, and the wavelength assigned to a coupler will automatically be assigned to all the transmitters and receivers that are members of the subnetwork represented by that coupler. To minimize the number of fibers, it is necessary (i) to assign $\min(W, C_T)$ and $\min(W, C_R)$ different wavelengths to each transmitting and receiving station, respectively, and (ii) to make equal use of each assigned wavelength at each station; i.e., the numbers of transmitters (receivers) in any given station which use two different wavelengths may differ by at most one. (It is also desirable, for standardization, to assign the same set of wavelengths to all stations.) We next present two wavelength assignment algorithms, along with combinations of (W, C_T, C_R) for which they achieve the maximum saving, namely a reduction by a factor of $\min(W, C_T)$ and $\min(W, C_R)$ in the number of fibers connected to transmitting stations and to receiving stations, respectively. Other combinations as well as correctness proofs for the assignment algorithms will not be presented.

Algorithm 1: arrange the couplers in groups of G, where G is the least common multiple of C_T and W; number the groups consecutively modulo W, beginning with 0. Next, number the couplers within each group consecutively modulo W, beginning with the group's number. The numbers correspond to distinct wavelengths. Applicable combinations:

- W divides C_T as well as C_R: each station is assigned all W wavelengths. Taken to the extreme, assuming $W{=}C_T{=}C_R{=}C$, a single fiber will be required from each station to the wiring closet.

- $W{=}C_T$ and $C_T > C_R$: a single fiber is connected to each station. Each transmitting station is assigned all W wavelengths, but different receiving stations may be assigned different subsets of wavelengths.

- $W > C_T$, $W > C_R$, W divides $C_T \cdot C_R$: a single fiber is connected to each station. Different stations may be assigned different subsets of wavelengths. Taken to the extreme $(W{=}C_T \cdot C_R)$, this is a pure WDM implementation.

Algorithm 2: number the couplers modulo W.The numbers again correspond to wavelengths. Applicable combinations:

- $W{=}C_R$, $C_R > C_T$: a single fiber is connected to each station. Each receiving station is assigned all W wavelengths, but different transmitting stations may have to be assigned different subsets of wavelengths.

The use of W different wavelengths also permits a reduction by a factor of W in the number of couplers, which is achieved by replacing each group of W couplers that carry different wavelengths by a single coupler. A grating could be used to combine the inputs of the individual single-wavelength couplers into multi-wavelength inputs of a single coupler. At the output of the coupler, they could be separated again using the same grating. Fig. 2 shows a possible implementation for the case $W{=}C_T{=}C_R{=}C$, wherein a single grating spectrometer is used to perform the wavelength multiplexing/demultiplexing; a representative $(\frac{N}{C} \times \frac{N}{C})$ coupler is shown, performing the mixing, splitting and interchange operations. The hardware requirements here would be similar to those for $\mathcal{PBI}$ when implemented with C wavelengths multiplexed onto a single broadcast network. While the reduction in the number of couplers is possible, it may not be economical due to the additional gratings and alignment stages that are required.

The combination of spatial and wavelength multiplexing can be used to support large values of C_T and C_R, while maintaining small cable size and manageable wavelength multiplexing.

It is very important to note that the wavelength–coupler tradeoff affects the number of couplers but has no effect on the size of each coupler, and that efficient merging of different wavelengths is possible. Consequently, the power budget is not affected. Also, the effect on hardware requirements is the same for $\mathcal{PBI}$ and for $\mathcal{SBI}$; consequently, the previous comparison is valid regardless of implementation.

Other multiplexing techniques. The same assignment algorithms that were used for WDM could be used with other multiplexing techniques, such as angular multiplexing,[12] modal multiplexing if and when it becomes possible.

Fiber-optic implementation using linear buses

In this case, $\mathcal{SBI}$ and $\mathcal{PBI}$ require the same size and number of couplers. In an SDM implementation, $\mathcal{SBI}$ would clearly C times more fibers, but if WDM is used the fiber requirements are also the same.

In the sequel, we focus on the centralized implementation.

Suitability to various access schemes

The operation of $\mathcal{SBI}$ in conjunction with access schemes that do not require a transmitting station to sense the channel over which it is transmitting is straightforward; ALOHA and TDMA[3] are examples of such schemes. However, operating it in conjunction with access schemes that do require such sensing (e.g. CSMA) is problematic, since the transmitting stations have no sensors and the signal to be sensed is not present at those stations. With $\mathcal{PBI}$, a simple solution is to merge each transmitting station with a receiving station and to use the receivers as sensors. This would work, provided that $N_T{=}N_R$. However, this cannot work for $\mathcal{SBI}$, since each transmitting station has only one subnetwork in common with each receiving station. Two ways of modifying $\mathcal{SBI}$ for use with such channel-sensing schemes are:

Modifying the stations. One could add to each transmitting station C_T sensors, which can often be much simpler than receivers. To make a signal that needs to be sensed available at the transmitting station, one could use a $(P \times (P+Q))$ coupler for each subnetwork and return an output of the coupler to each of the P relevant transmitting stations; this would double the number of fibers connected to each transmitting station. Alternatively, one could use a $(P \times (Q+1))$ coupler with a mirror on one of the outputs, causing some of the power to return to the transmitters. The fibers from the transmitting station to the wiring closet will thus be used in both directions, and taps would be required to enable the sensors to sense the reflected signals.

Modifying the interconnection. Letting $N_T = N_R$ and $C_T = C_R = C$, we again note that the number of groups of transmitting stations and the number of stations in each group are equal to the corresponding numbers for receiving stations. Next, let us think of the nth transmitting station and of the nth receiving station as two halves of the nth bidirectional station. Observing that for each (i, j) such that $i \neq j$, there is one subnetwork for transmissions from the ith group of stations to the jth group of stations and a different subnetwork for transmissions from the jth group to the ith group, we merge each such pair of subnetworks into a single subnetwork. Each member station can both transmit and receive over such a combined subnetwork; such a subnetwork is therefore a conventional broadcast network for the stations that belong to groups i and j, and can therefore be implemented with any access scheme. Finally, since intragroup communication can take place over any of the $(C-1)$ combined subnetworks of which a group's stations are members, there is no need for dedicated intragroup subnetworks. We can therefore delete those subnetworks $(i=j)$, and save one transmitter and one receiver per station. The end result is that, for $(C-1)$ transmitters and receivers per station, we have $C(C-1)/2$ subnetworks. Therefore, for stations with C transmitters and receivers, we can construct $C(C+1)/2$ independent subnetworks, each of which serves the members of two groups and can be implemented with any access scheme. The fiber-optic implementation was embedded in the above description; the required couplers would be of size $(\frac{2N}{C} \times \frac{2N}{C})$.

Finally, we wish to present an interesting scheme for operating SBI efficiently in the case that $C_T = C_R = C$ and $N_T = N_R$. In this case, the groups of transmitting stations and receiving stations are of equal size. One could therefore merge each transmitting station with a receiving station to form C groups of bidirectional stations. It can readily be seen that there are C intragroup subnetworks and $C(C-1)$ intergroup ones. Each intragroup subnetwork could be used by the group's members for the scheduling of their transmissions over the remaining subnetworks of which they are (the only) transmitting members. This subnetwork would also be used for intragroup communication.

Practical Considerations

Saving transmitters . SBI appears to require C_T transmitters per station. However, when implemented with SDM, it is possible to implement a transmitting station using fewer transmitters. This requires a switch at the station, which can steer the transmitters' outputs to any of the C_T subnetworks over which the station has to transmit. As long as the station's average data rate is low (less than B times the number of actual transmitters), its throughput will not be affected.

In practice, particularly if inexpensive LED's are used for transmission, it may be desirable to have a transmitting element connected to each outgoing fiber and to switch electrical signals between those transmitters, rather than to switch optical signals. By doing so, the optical subsystem remains completely non-switched. In conjunction with this, it may be noted that arrays of 12 individually addressable LED's, directly coupled to multi-fiber ribbon cable, have been developed.[13]

It should be stressed that switching at the transmitting station operates without tying up the network and may therefore be very slow. It is also worth noting that the decision whether or not to use fewer than C_T transmitters can be made independently for each station, based on its expected traffic load, thus allowing a significant cost reduction for small users without restricting the large ones. Finally, we note that transmitters can also be saved in PBI if it is used for bit-serial transmission, but not if it is used in parallel for the transmission of C-bit words.

Saving receivers. The situation for receivers is more complicated, since one does not know on which subnetwork to expect the incoming packet (unless TDMA or reservation schemes are used to implement the subnetworks). Therefore, one would need address-detection capability on all C_R incoming fibers, as well as a fast switch.

Additional issues .

- The $(P \times Q)$ couplers can be $(\max(P,Q) \times \max(P,Q))$ optical stars (not all fibers of which are used if $P \neq Q$), or networks made from elementary couplers. The particular solution adopted will depend on the power margin, on the cost of couplers, etc.

- At the present time, a maximum of five different wavelengths have been multiplexed onto a single fiber from an integrated laser array,[14] and ten from separate lasers,[15] and it may be that wavelength multiplexing alone will never achieve desired values of C. Similarly, if one wishes to use a single multi-fiber cable for each station, or, eventually, multi-core single fibers,[16] there is a limit on the degree of spatial multiplexing. In such cases, a hybrid implementation can be used.
- Since optical transmitters are currently fixed-wavelength sources, each transmitting station must have at least one optical transmitter for each of the wavelengths it uses.
- Factors that are to be taken into consideration when choosing C_T and C_R, given their product: $S \propto (C_T \cdot C_R)$; number of fibers $\propto (C_T + C_R)$ for SDM implementations; cost of receivers $\propto C_R$; cost of transmitters $\propto C_T$, unless transmitting stations are implemented with fewer than C_T transmitters; power is split $\max(Q, P)$ ways if reciprocal couplers are used; since the optimal power splitting factor is Q, one should make sure that $Q \geq P$ if power budget is critical.

Extensions

In this section, we briefly explore several extensions of the basic SBI that has been discussed thus far. A more comprehensive discussion of the underlying theory will be presented elsewhere.

Guaranteed concurrency and fault tolerance

As we have already mentioned, the performance of SBI is more sensitive than that of PBI to non-uniformity of the traffic requirement, since each source-destination pair is connected by only one subnetwork, and packets of source-destination pairs that communicate through the same subnetwork cannot be accommodated concurrently. For the same reason, any single failure breaks the full connectivity. We next present ways of trading maximum concurrency and power for guaranteed minimum concurrency and fault tolerance (maximum and minimum are over all possible traffic requirements).

Multiple paths between stations. There are several ways to construct multiple paths between each pair of stations:

- Building k independent SBI's, each using C_T/k transmitters per station and C_R/k receivers per station. This results in a maximum concurrency of $C_T \cdot C_R/k$, with a guaranteed concurrency of at least k (potentially higher if different assignments of stations to groups are used in the different SBI's and several conditions on the values of C_T, C_R, k are met).
- Building k SBI's, each utilizing all the transmitters and receivers, such that each SBI has its own couplers, but they are connected at the stations. (Whenever a transmitter transmits, it does so over one subnetwork of each of the k SBI's.) Obviously, the stations must be arranged differently in each SBI.
- Combining PBI of size k with SBI which uses the remaining transmitters and receivers. Maximum concurrency is $k + (C_T - k)(C_R - k)$, and the guaranteed concurrency is $k + 1$. If, for example, C_T=C_R=C and k=$C/2$, maximum concurrency is $(\frac{C^2}{4} + \frac{C}{2})$, which is on the order of C^2, and guaranteed concurrency is $(\frac{C}{2} + 1)$, which is on the order of C. Fig. 3 illustrates this example.

2-hop transmission with randomization. Using 2-hop routing, wherein the first hop is from the station that originates the packet to a station that is chosen at random, and the second hop is from the latter to the true destination of that packet, guarantees a concurrency of C and reduces the potential concurrency to $C^2/2$. It can also be shown that additional hops cannot help. While this scheme is somewhat more efficient than multiple paths, it is brought here only as an anecdote since it is not a single-hop scheme.

In all the above schemes, the number of faults that can be tolerated (broken fibers, faulty transmitters, faulty receivers) is equal to the guaranteed concurrency.

Asymmetric topologies

It can readily be seen that one can distribute the stations (transmitting as well as receiving) unevenly among the groups, as long as the number of groups is unchanged. Coupler sizes would have to be adjusted appropriately, but the potential concurrency remains unchanged. Also, with SDM implementations, a transmitting station that has fewer than C_T outgoing fibers can have some of its fibers connected to more than one subnetwork, while preserving the optimality of the power budget. The same could also be done for stations with fewer than C_R incoming fibers, but in this case power optimality might be lost if reciprocal couplers are used. Asymmetric topologies may be useful in the case of a nonuniform traffic pattern, or when constraints such as geographical location of stations make them desirable.

More general topologies

The discussion has so far been limited to a restricted class of SBI. This class is characterized by the fact that two transmitting stations (likewise two receiving stations) are either connected to the exact same subnetworks or have no common subnetwork. In general, this need not be the case. The general rule for generating all single-path SBI's is simply to use one of the $C_T \cdot C_R$ connection possibilities for each source-destination pair. The topologies generated by this rule will all preserve power optimality in the sense of the minimal power split, as long as each transmitter of a station is used for the same number of connections. However, if one uses reciprocal couplers, true power optimality can only be achieved if, in addition to the above, the nodes of the directed graph corresponding to the physical topology that achieves the interconnection (these nodes are transmitters, couplers and receivers; the edges are fibers), all have an indegree smaller or equal to their outdegree, with the exception of receivers which may have indegree one. The example of SBI presented in Fig. 1 clearly meets the requirements for power optimality.

Applications

SBI is not presented merely as an architecture to replace existing broadcast LAN's, but rather as an approach capable of extending the usefulness of such LAN's to accommodate anticipated future applications requiring higher throughput than can be provided today.

Consider, for example, the design of a campus-wide browsing library, where users would scan rapidly through pictorial information sent by servers located in the libraries. Several hundred individuals might want to use the system at any one time, and each one of them might require several megabits per second of bandwidth (average); the corresponding aggregate throughput could reach into the gigabit per second range, well beyond the current capability of any commercially available LAN. SBI, however, could provide an adequate solution. Note that each individual terminal need not be equipped with multiple transceivers: a number of such terminals could share a small conventional LAN (serving a floor of a building, a department, etc.); each such LAN would have a gateway to a backbone SBI, that would interconnect those LAN's and the file servers.

SBI could also be useful in future medical environments, where it is anticipated that high-resolution images obtained by a variety of means, such as CAT scanners, MRI, ultrasound, PET, standard X-rays, etc., will be archived in a number of locations, and retrieved on graphics display units located in wards, operating rooms, etc., in order to facilitate diagnosis, follow patient evolution, etc. The aggregate throughput required in such an environment would make it a suitable candidate for SBI.

A cost-effective design of SBI for such applications will be facilitated by the fact that SBI can be implemented with different numbers of transmitting stations and receiving stations, a station need not have the same number of transmitters and receivers, and the actual number of transmitters in a transmitting station can be smaller than the number of subnetworks of which it is a member.

Conclusion

We have introduced the selective broadcast interconnection for local area networks. By employing many subnetworks, each of which interconnects only a subset of the stations, SBI has concurrency and aggregate throughput which, for uniform traffic requirements, increase quadratically with the number of transmitters and receivers per station. Viewed differently, any given throughput can be achieved with fewer and/or slower transmitters and receivers per station than would be required with parallel broadcast networks, each interconnecting all stations. The cost/performance ratio of SBI should be more favorable than that of the more conventional topologies, particularly for large network sizes. SBI is also superior in terms of power budget. Another attractive feature of SBI is that it does not use switches; being entirely passive, it should be reliable and maintenance-free. SBI could find applications wherever a moderate-to-large number of stations need to communicate at very high rates.

We noted that generalizations of the basic SBI can serve to trade maximum concurrency and power for guaranteed concurrency and fault-tolerance, and that fault-tolerance can also be increased if high-level protocols provide for message-forwarding in the event of partial network failure.

We feel that SBI is uniquely suited for optical fibers, the small size of which makes spatial multiplexing an attractive possibility. Nevertheless, it can also be used with other technologies, such as broadband LAN's and packet radio networks. (In both cases, frequency-division multiplexing would most likely be used.)

Acknowledgments

This work was supported by the National Aeronautics and Space Administration under grants NASA-NAG-2-292 and NASA-NAGW-419. Y. Birk is also supported by an IBM graduate fellowship.

References

1. Finley, M. R., IEEE Communications Magazine, Vol. 22, No. 8, p. 22. 1984.
2. Bergman, B. L., in Fiber Optics and Buyers Guide (Information Gatekeepers, Boston), p. 22. 1984.
3. Tobagi, F.A., in Computer Network Architectures and Protocols, P.E. Green, Ed. (Plenum press, New York, NY), Chap. 6, pp. 145-189. 1982.
4. Tobagi, F. A., Borgonovo, F. and Fratta, L., IEEE J. Select. Areas Commun., Vol. JSAC-1, No. 5, Nov. 1983.
5. Fine, M. and Tobagi, F.A., IEEE Transactions on Computers, Vol. C-33, No. 12, pp. 1130-1159. Dec. 1984.
6. Muoi, T. V., J. of Lightwave Technology, Vol. 2, p. 243. 1984.
7. Villaruel, C. A., Moeller , R. P., and Burns, W. K., IEEE J. Quantum Electron., Vol. 17, p. 941. 1981.
8. Eng, S.T., Tell, R., Andersson, T. and Eng, B., J. of Lightwave Technology, Vol. 3, No. 4. Aug. 1985.
9. Hara, E. H., in Proceedings of Global Telecommunications Conference (IEEE, Piscataway, N.J.), paper E1.4, p. 961. 1982.
10. Saniter, J., Schmidt, F. and Werner, W., in Digest of Topical Meeting on Optical Fiber Communication (Optical Society of America, Washington, D.C.), paper TUJ3, p. 44. 1985.
11. Marhic, M.E., Opt. Lett., Vol. 9, p. 368. 1984.
12. Stearns, R.C., Asawa, C.K. and Yao, S.K., J. of Lightwave Technology Vol. 2, p. 358. 1984.
13. Deimel, P.P. et al, in Digest of Topical Meeting on Optical Fiber Communication (Optical Society of America, Washington, D.C.), paper TUC4, p. 32. 1985.
14. in Lasers and Applications, Vol. 4, p. 18. January 1985.
15. Olsson, N.A., Hegarty, J., Logan, R.A., Johnson, L.F., Walker, K.L. and Cohen, L.G., in Digest of Topical Meeting on Optical Fiber Communication (Optical Society of America, Washington, D.C.), paper WB6, p. 88. 1985.
16. Sumida, S., Maekawa, E. and Murata, H., J. of Lightwave Technology, Vol. 3, p. 159. 1985.

Selective broadcast interconnection (*SBI*)

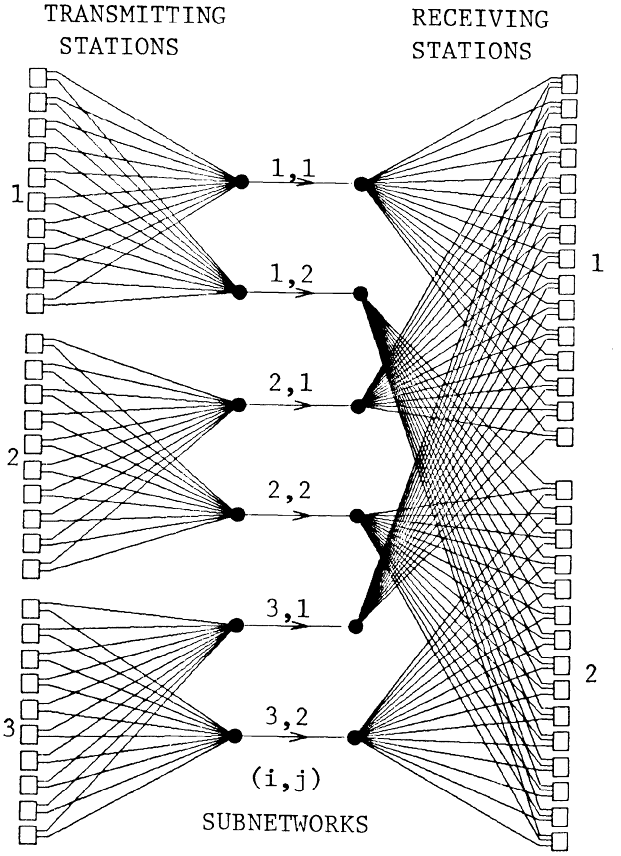

Fig. 1 Connection diagram of *SBI*. Rectangles represent individual stations, and lines ending at the stations represent individual transmitters and receivers. $N_T=N_R=30$, $C_T=2$, $C_R=3$; $P=10$, $Q=15$.

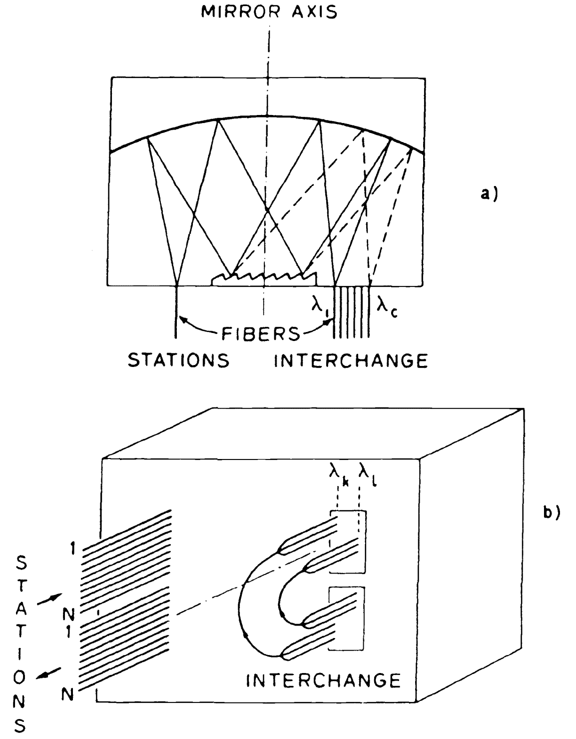

Fig. 2 Implementation using WDM. a) Top view of the internal arrangement of the grating spectrometer; b) Perspective view of the outside of the spectrometer, showing the fiber layout and two representative $(\frac{N}{C}\times\frac{N}{C})$ couplers.

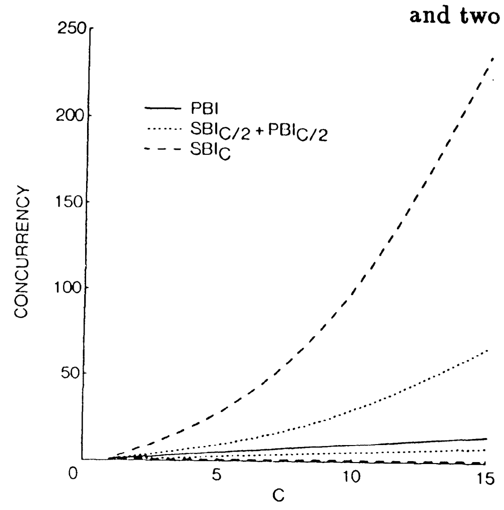

Fig. 3 Minimum and maximum effective concurrency for *PBI*, for $[SBI_{C/2}+PBI_{C/2}]$ and for *SBI*. $C_T=C_R=C$.