Selective broadcast interconnection: a novel scheme for fiber-optic local-area networks

M. E. Marhic,* Y. Birk, and F. A. Tobagi

Computer Systems Laboratory, Department of Electrical Engineering, Stanford University, Stanford, California 94305

Received April 29, 1985; accepted September 18, 1985

We introduce a passive, unswitched scheme for directly interconnecting N stations, each of which has C transmitters and receivers. Implementations using fiber optics with spatial multiplexing and optionally wavelength multiplexing are discussed. This scheme utilizes the same resources as standard topologies with C parallel buses but outperforms them in two respects: (1) the aggregate throughput is proportional to C^2 rather than to C and (2) the power of each transmitter need reach only N/C, instead of N, receivers.

Many configurations considered to date for fiber-optic local area networks (FOLAN's) are of the conventional broadcast type, in which all N stations use time-division multiplexing (TDM) to share a common communication channel (bus).[1,2] Each station has a single transmitter and a single receiver, both of peak data rate B. Only one successful transmission can take place at any one time; consequently the aggregate throughput S cannot exceed B, and the average data rate per station, D, is at most B/N. For large N and/or D, this forces the transmitters to work at high speed during short intervals, which is undesirable. Also, each receiver spends most of its time listening to communications intended for others, which is generally a waste of resources. These undesirable features of TDM make the design of broadcast FOLAN's more difficult as N and/or D is increased.

Furthermore, conventional broadcast FOLAN's are subjected to a power limitation when implemented with passive 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. 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,[3] the maximum N that can be accommodated by a broadcast FOLAN for a given P_T decreases with increasing B and can become quite small at high speeds; this is especially true if a simple but energy-inefficient structure such as a linear bus is used.[4]

Most other FOLAN's considered to date are of the ring or loop type, which employ point-to-point optical links and are subject to restrictions on B and S 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 conditions). Proposals have been made to extend the capabilities of rings into the gigabit-per-second range,[5,6] but the severe limitations imposed by TDM may greatly hamper these schemes.

In order to increase S beyond B or $2B$, one must use systems with a high degree of concurrency. To achieve this with point-to-point optical links, one could use intelligent nodes (some or all of which may be the stations) that route messages according to their destinations. Several such methods have been proposed for FOLAN's,[7,8] and many more could be adapted to them when fast optical and/or electronic switching and electronic processing become commonplace.

For broadcast networks, a straightforward approach to achieving concurrency is to transmit the messages over any of C broadcast subnetworks, each connecting all stations; these subnetworks could be used either independently for bit-serial transmission or for the transmission of C-bit words. In either case, the hardware required at the stations is comparable with that of C independent broadcast networks; an implementation using C parallel fibers everywhere (spatial division multiplexing, or SDM) would also require replication of the interconnection hardware (fibers and couplers), but the use of wavelength-division multiplexing (WDM) could obviate that requirement. In any case, S is proportional to the amount of communication hardware at the stations (C). We denote this approach PBI (for parallel broadcast interconnection) and use it as a reference in the evaluation of our scheme.

In this Letter, we propose a novel method to achieve high concurrency in broadcast networks, the selective broadcast interconnection (SBI) scheme. In SBI, as in PBI, there are broadcast subnetworks and multiple transceivers per station. The basic difference is that, unlike in PBI, each subnetwork in SBI connects only a subset of the transmitting stations to a subset of the receiving stations, in such a way as to have one and only one transmission path between each transmitting station (TS) and each receiving station (RS). Whenever a station wants to send a message to some other station, it uses the transmitter that is on the subnetwork connecting the two stations. Using the same amount of communication hardware at the stations as in PBI, SBI permits a concurrency of the order of C^2 rather than of C. For the sake of brevity, we present here a simple description of a particular case of SBI; more general possibilities will be presented elsewhere.

0146-9592/85/120629-03$2.00/0

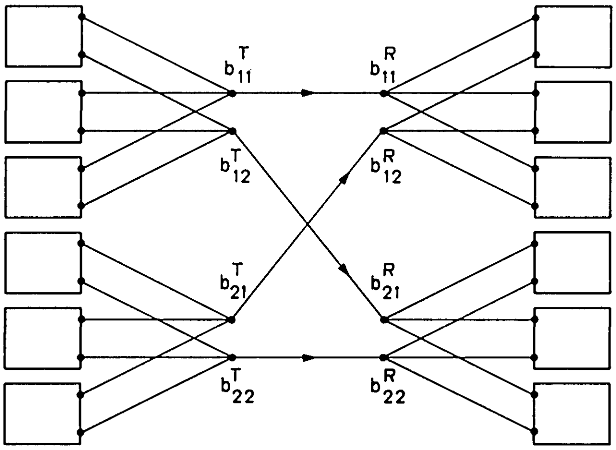

Fig. 1. Schematic of SBI for $N = 6$, $C = 2$, $Q = 3$. The rectangles represent the transmitting stations and the receiving stations. The individual transmitters and receivers are not shown, but the lines ending at the stations are connected to them.

For $N = QC$, where Q and C are integers, SBI can be constructed as shown in Fig. 1. For simplicity, the TS's and the RS's are shown at opposite ends of the diagram, even though they might in practice coincide in space. The lines represent an abstract schematic of the connections, but they can actually be interpreted as representing optical fibers and thus as an SDM implementation. The SBI of Fig. 1 is formed as follows. We arrange the TS's in C disjoint groups of Q stations each. Next, within each such group, we form C disjoint bundles of Q fibers each, such that the jth bundle contains the jth fiber of each TS in the group. We denote by $b_{ij}{}^T$ the jth bundle of the ith group. There are C^2 such bundles. We then perform the same operation with the receivers, also forming C^2 bundles, denoted $b_{ij}{}^R$. Finally, we optically couple all fibers of $b_{ij}{}^T$ to those of $b_{ji}{}^R$ by means of a broadcast $Q \times Q$ directional coupling network (hereafter referred to as a $Q \times Q$ coupler). There are C^2 such couplers altogether, providing a single optical path between each TS and each RS. Consequently, the degree of concurrency in this arrangement can potentially reach the value C^2, compared with C for PBI.

The other important way in which SBI outperforms PBI is in terms of power splitting: The power of each transmitter need only reach Q receivers in SBI, compared with N receivers in PBI. In fact, no other scheme for directly interconnecting N stations with C transceivers each can be more power efficient than SBI, since the power of each transmitter must be split at least $N/C = Q$ ways in order to achieve connectivity. This power gain of C over PBI can be beneficial in helping to increase passive network sizes or in reducing the cost of the transceivers.

So far, we have interpreted Fig. 1 as an actual fiber-optic implementation using SDM. This is indeed feasible in practice and, provided that C is not too large, can be realized by using single cables containing multiple fibers, which are relatively easy to handle and install. Multicore single fibers[9] might also eventually be useful for SDM implementation. Each $Q \times Q$ coupling network can be implemented by using all the approaches developed for conventional broadcast networks. If a wiring closet is used, then SBI uses the same amount of fiber as a similarly configured PBI. The latter, however, uses only C $N \times N$ couplers, as compared with C^2 $Q \times Q$ couplers for SBI, and so the PBI might have an advantage on this basis. Qualitative comparisons based on costs of large star couplers indicate that these do not give PBI a substantial advantage over SBI. A quantitative comparison may be obtained by assuming that all couplers under consideration are to be made from elementary 2×2 couplers.[10] Let us assume further that $N = C^2$ and that C is a power of 2. Under these circumstances, it can be shown that SBI requires only half as many 2×2 couplers as PBI.

Should one find the large fiber counts of SDM objectionable, one could implement the C point-to-point links between each of the stations and a centrally located SBI by means of WDM. In that case, the simplest thing to do is to use a separate transmitter for each wavelength at each transmitting station, multiplexed onto just one fiber. N such fibers would be brought to a central location, where demultiplexing would take place, yielding the NC possible signals from all stations. These could be fed into individual fibers, and then treated exactly as in the case of SDM, by means of the same $Q \times Q$ couplers. The interchanged signals would then reach the RS's on N other fibers and be demultiplexed. Here one should make sure that each coupler receives signals at only one wavelength and that each RS receives each wavelength

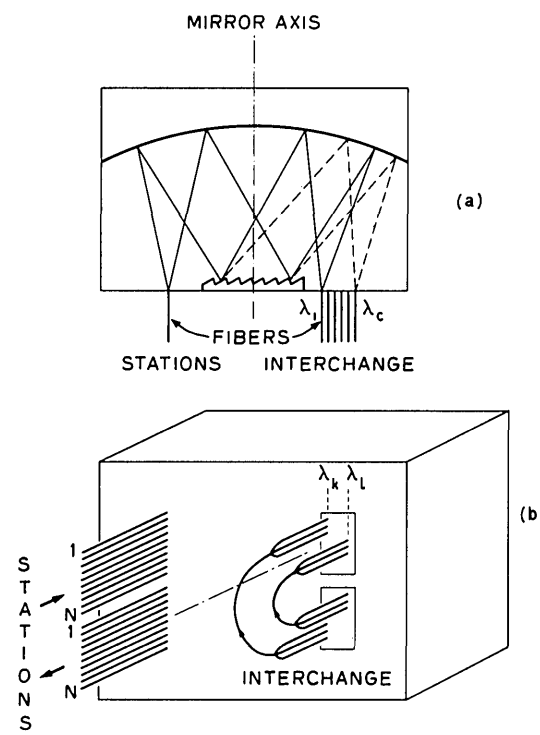

Fig. 2. Implementation using WDM. (a) Top view of the internal arrangement of the grating spectrometer; (b) perspective view of the outside of the spectrometer, showing the fiber layout and two representative $Q \times Q$ couplers.

once and only once. Figure 2 shows a possible implementation of SBI by WDM, wherein a single grating spectrometer is used to perform the central multiplexing/demultiplexing; a representative $Q \times Q$ coupler is shown, performing the mixing, splitting, and interchange operations.

At the present time, five different wavelengths from an integrated laser array,[11] and ten from separate lasers,[12] have been multiplexed onto a single fiber. If these numbers of wavelengths are insufficient to implement SBI by WDM alone, WDM can be combined with SDM. In this hybrid situation, the interchange will take on a structure that is still physically realizable. One might also consider using angular multiplexing[13] if this technique ever becomes practical.

Another practical advantage of SBI is that the structure of the TS's could be substantially simplified by using switches: in the case of SDM, integrated-optic or fiber-optic switches could route the output of a single light source to each of several fibers. Cross talk, response time, insertion loss, and cost would have to be considered in assessing the benefits of such a modification. Note that this is not an option in the case of PBI with C-bit words. No similar saving in the number of optical receivers is possible, because switching into a single such receiver would require knowledge of arrival time and origin of messages, which is generally not available (unless a deterministic time slot access scheme is used).

In conclusion, we have introduced a novel, switchless, passive interconnection method to achieve concurrency of the order of C^2 in broadcast networks with C transceivers per station. The method circumvents the disadvantages of TDM by utilizing beneficial aspects of SDM and/or WDM, and it makes effective use of the available communication resources, from the points of view of both information rate and power. With moderate values of C, SBI should exhibit aggregate throughputs in excess of what can be achieved with conventional broadcast FOLAN's and straightforward extensions thereof. Although SBI could clearly be implemented by means of other technologies, it is particularly well suited to fiber optics because of their small size and anticipated low cost, which should make SDM attractive. SBI could find applications wherever a moderate to large number of stations need to communicate at a high aggregate rate, such as in browsing through on-line libraries and archives.

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

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

References

1. M. R. Finley, IEEE Commun. Mag. **22**(8), 22 (1984).
2. B. L. Bergman, in *Fiber Optics and Buyers Guide* (Information Gatekeepers, Boston, Mass., 1984), p. 22.
3. T. V. Muoi, J. Lightwave Technol. **LT-2,** 243 (1984).
4. C. A. Villaruel, R. P. Moeller, and W. K. Burns, IEEE J. Quantum Electron. **QE-17,** 941 (1981).
5. H. F. Taylor, IEEE J. Selected Areas Commun. **SAC-1,** 500 (1984).
6. S. D. Personick, in *Digest of Topical Meeting on Optical Fiber Communication* (Optical Society of America, Washington, D.C., 1984), paper TUJ1, p. 44.
7. E. H. Hara, in *Proceedings of Global Telecommunications Conference* (Institute of Electrical and Electronics Engineers, 1982), paper E1.4, p. 961.
8. J. Saniter, F. Schmidt, and W. Werner, in *Digest of Topical Meeting on Optical Fiber Communication* (Optical Society of America, Washington, D.C., 1985), paper TUJ3, p. 44.
9. S. Sumida, E. Maekawa, and H. Murata, J. Lightwave Technol. **LT-3,** 159 (1985).
10. M. E. Marhic, Opt. Lett. **9,** 368 (1984).
11. Lasers Appl. **IV,** 18 (1985).
12. N. A. Olsson, J. Hegarty, R. A. Logan, L. F. Johnson, K. L. Walker, and L. G. Cohen, in *Digest of Topical Meeting on Optical Fiber Communication* (Optical Society of America, Washington, D.C., 1984), paper WB6, p. 88.
13. R. C. Stearns, C. K. Asawa, and S.-K. Yao, J. Lightwave Technol. **LT-2,** 358 (1984).
14. W. Stallings, *Local Networks* (Macmillan, New York, 1984), Chap. 5.