B PHYSICS...DIGITAL TRIGGERS & DATA ACQUISITION USING NEW MICROPLEX & DATA COMPACTION ICs, MANY PARALLEL MATH PROCESSORS, & FIBER OPTICS TO FAST DIGITAL TRIGGERS & HIGHER LEVEL FARMS

E. Barsotti, M. Bowden, H. Gonzalez and C. Swoboda
Fermilab
Batavia, Illinois USA

Summary

This paper gives a systems engineering approach to solving the triggering and data acquisition problems in a B physics experiment. A strong emphasis is put both on using standard busses, data transmission techniques, etc. wherever possible and on minimizing 'special' electronics built with combinatorial and, at best, minimally programmable logic. The increased use of buffering and parallel processing using high-level language programmable processors at various levels of the experiment is strongly recommended.

The following notes give suggestions for:

1. A **new microplex-like IC** containing FADCs and scanning and buffering logic.
2. A **new data compaction IC** which orders good channels (i.e., addresses and, if present, digitized FADC data) into blocks of memory to be fed into a digital fast trigger array of math processors and subsequently into higher level processing farms.
3. A **digital few-microsecond trigger system using an array of math processors.**
4. **High-rate data transmission** into both the fast trigger processor array and a higher level farm **using fiber-optic cables**...a technology which surpasses copper cables in performance already even though it is in its infancy.
5. **International standards efforts for** very high-speed point to point and point to multiple points **data transmission via fiber-optic cables.**
6. Present and near future **processing capabilities in high level farms.**

The Problem

Figure 1 is a very simple description of the triggering and data acquisition problem in a B physics collider experiment. Note the several gigabyte per second (GB/s) data rate into the fast trigger math processor array or an even higher data rate from the 'new microplex-like ICs' into the first column of new data compaction ICs. Note also that we are using highly parallel digital, not analog, methods for fast triggering. Math processors can do floating point Ax + B calculations in 30 to 60 nanoseconds. Figure 1 assumes a factor of five reduction by a fast analog Level 1 trigger.

New Readout & Data Compaction ICs

Figure 2 shows how new microplex-like ICs and new data compaction ICs are used to form a pipelined multiple stage readout and data compaction subsystem capable of pipelining at rates of a few hundred nanoseconds per pipeline stage. Note that each pipeline stage reduces the maximum stored data by a factor of two. It is estimated that less than one tenth of the

160,000 channels of data are hit. The output of this logic is several data compaction ICs each containing ordered sets of good channel data (i.e., addresses and, if present, digitized FADC data). This data is used as input to the fast trigger math processor array and, if accepted by the fast trigger, as eventual input to the high level processing farm.

Each new microplex-like IC has two identical sections, each having one FADC, as shown in Figure 3 which scan analog-compared inputs internally at 200MHz or 500MHz rates. When a good hit is detected, an FADC conversion is done. For good hits the scan rate is slower. We have been told FADCs can be integrated into digital ICs. We do not know details at this time. It is possible that this one new design can be used for both yes/no only data and data requiring a digitized analog value.

Each new data compaction IC also has two identical sections and, in fact, has much the same circuitry as the new microplex-like IC as shown in Figure 3. The scanning rate of each section need only be of the order of 40MHz. Ordered sets of addresses and, if present, FADC data are the inputs and outputs of these ICs. Note that each data compaction ICs can only store one half of its possible input number of data channels. We thought this might be acceptable given that less than one tenth of the inputs result in hit channels. Modifications to this factor of two scheme should be straightforward. Here again we need additional input from physicists.

Fiber-Optic Data Transmission

The FASTBUS Standards Committee has begun work specifying a standard fiber-optic high-speed data transmission point to point (or to multiple point) link. Similar efforts are starting in ESONE, a European standards body, and within the IEEE. Communication between these organizations will take place hopefully leading to one standard.

The transmission of data from front-end data acquisition systems to fast trigger math processor arrays or processor farms using fiber-optic cable is possible at rates exceeding 100 MBytes /sec. At 100 MBytes/sec the bit transmission rate on the fiber must be approximately one Gigabit/sec. Efficient use of the fiber bandwidth therefore requires that data be applied to the optical transmission system at approximately one byte every 10 ns. The feasibility of feeding and removing data from an optical system at this rate is possible but will be costly. Today a 1GHz optical link alone costs approximately $12,000. The development costs of high-speed buffering, data encoding, data transmission, and data recovery must be added to the optical component costs.

A more realistic approach would be to use VLSI chips presently available to format and transmit data over fiber-optic cable. As the technology advances, the commercial development of VLSI chips that allow higher data rates to be injected and removed from the optical link will become available. As these devices become available, the implementation costs will drop.

This approach is presently under consideration by the FASTBUS Standards Committee and by others, both in the United States and Europe. In particular, the use of AMD Supernet chips to implement 12.5 MByte optical links is planned. These chips are available today. They provide data encoding, data serialization, data recovery and limited error checking. When connected to optical transmitters and receivers, these chips provide transparent point-to-point data transmission links.

Figure 4 depicts single and multiple fiber, point-to-point links using the AMD Supernet chips which could be used to transmit data from 12.5 to 50 MBytes/sec. Higher data transmission rates can be accomodated by increasing the number of fiber links. The application of this technology to an experiment would take the form shown at the bottom of Figure 4.

Data from FASTBUS, CAMAC, VME, and detector mounted electronics could be transmitted over standardized optical links to memory modules. Present technology would allow the fiberoptic cable to be daisy-chained to additional memory modules thereby allowing data from a single front-end system to be transmitted to multiple memory modules. This data could then be read by processor farms for further analysis.

The implementation of the data transmitter is shown in Figure 5. Parallel data from a readout device is fed in parallel to the optical transmitter. The data is transmitted serially to a module implemented as shown in Figure 6. This module accepts data from four different sources or 4 fibers from a single device and multiplexes 32 bit data words onto a FASTBUS backplane. The transmission of data from front-end equipment over fiber-optic cable provides a high bandwidth media for future data transmission rate upgrades and excellent electrical isolation between the detector and the data acquisition system.

Data Routing Into Level 3

A major goal in a high throughput data acquisition system is to maintain parallel data paths at all levels. The Event Building function in particular, must be implemented in parallel. This can easily be accomplished, as shown in figure 7, with an array of dual-ported event memories. Scattered data from the preceding level is transmitted in parallel over standard point to point serial links to a single row of event memories. A processor in that row builds the event by reading the event memories sequentially. By balancing the number of processor rows with input data links, the full input bandwidth is preserved. The number of processors in each row depends on the average event processing time.

With current technology fiber-optic encoders and standard busses, a throughput of 40 MBytes/sec per data link is typical. A system with a memory array of 20 X 20 modules (1 crate by 20 rows) can therefore move approximately 800 MBytes/sec. A rate of only 80 MBytes/sec is initially required for the B physics data acquisition system.

Data Routing Into Level 2

While the dual-ported event memory architecture is very flexible, it suffers from two limitations which make it less feasible for Level 2 data rates. First, the memory array size can expand in proportion to the square of the data rate, which limits the practical system size to a less than 1 GByte/sec. Secondly, the use of a standard backplane in the horizontal datapath limits processor I/O. Both of these problems can be partially offset in a higher level processing system by the arrangement in Figure 8. In the case of Level 2 however, this is still insufficient. A FASTBUS crate can hold as many as 200 processors, resulting in I/O contention if the average processing time is less than 16 msec per 3.2 KByte sub-event (assumes 10% of the nominal event size used in the Level 2 trigger) Additional processor crates in the horizontal datapath only worsen the situation. For simplicity in calculation, assume a Level 2 event processing time of 512 μsec. Each processor then takes data at a rate of 6.25 MBytes/sec. Using current technology, this rate can be supplied using 1 fiber-optic cable per VME processor board or 4 cables per FASTBUS board. In the future, one fiber-optic cable per FASTBUS board is expected.

This alternate architecture (for data rates above 1 Gbyte/sec) is shown in Figure 9. All data transmission is via serial links in a packet switching network. In the simplest approach, the network is nothing more than a parallel N-bit circular shifter. Between each packet time slot, the outputs are rotated one position to perform the Event Building function. Multiple events must be buffered in the sending and receiving processors in order to fully utilize the network. With 100 Mbps links, a throughput of 1.6 GBytes/sec requires 128 serial channels

which is comparable to 40 standard parallel backplanes. Since multiple serial links can be routed to each processor crate or board, the backplane bandwidth restriction is eliminated. This approach requires synchronization of data packets between processing levels in contrast to the dual port memory architecture which is loosely coupled through a Data Flow Control Network (standard LAN).

A more flexible packet switch (used in telephone data switching networks) can be built using bitonic sorting elements. This network is self-routing but requires clocking of the node control logic which reduces the data rate to less than 50 Mbps per channel. It is also less deterministic.

Math Processor Arrays and Processor Farms

Several single chip Digital Signal Processors optimized for high throughput floating point calculations will be introduced in 1988, with performance in the 7-15 MW (MegaWhetstone) range. By comparison, the VAX 780 and first generation ACP modules are 1 MW machines. This means that performance levels of at least 15 VAX equivalents on a VME board and 50-60 VAX equivalents on a FASTBUS board are possible at a cost of about $200 per MW. Performance should double again within the time-frame of this detector.

System performance versus number of crates is shown in Figure 7 for the Level 3 architecture. Processing power would remain the same, but at higher potential data rates, in a fully serial packet network of similar size.

Original estimates for Level 3 data rates are modest (< 80 MBytes/sec) and can be handled by a system with 2 or 3 processor rows. Processing times as high as 1 second per event have been suggested (1000 events/sec = 1000 processors = 6 FASTBUS crates = 30 VME crates = $2 million).

The data rate into Level 2 is substantially higher, between 1 and 2 GBytes/sec. Assuming a half million events/sec and ~500 μsec/event requires 250 processors. A specialized Level 2 processor is unlikely to provide any cost or performance advantages. 500 μsec of processor time is sufficient for approximately 5-10,000 floating point operations.

B Physics.....The Problem

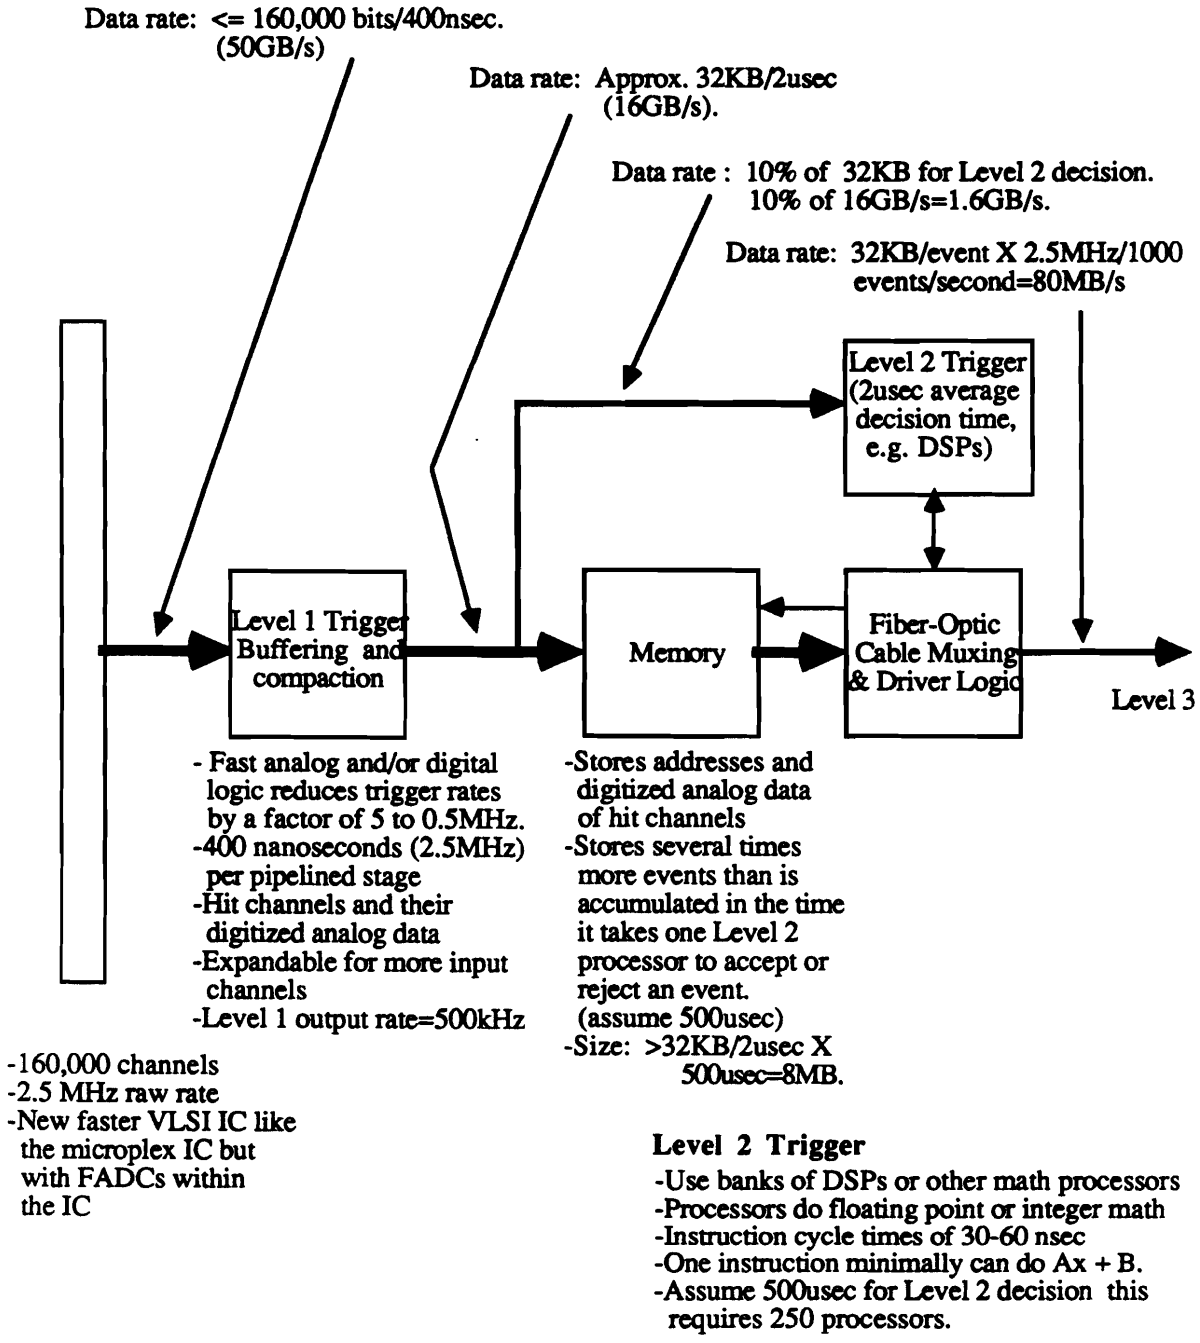

Figure 1

B Physics...Front-End Trigger Logic

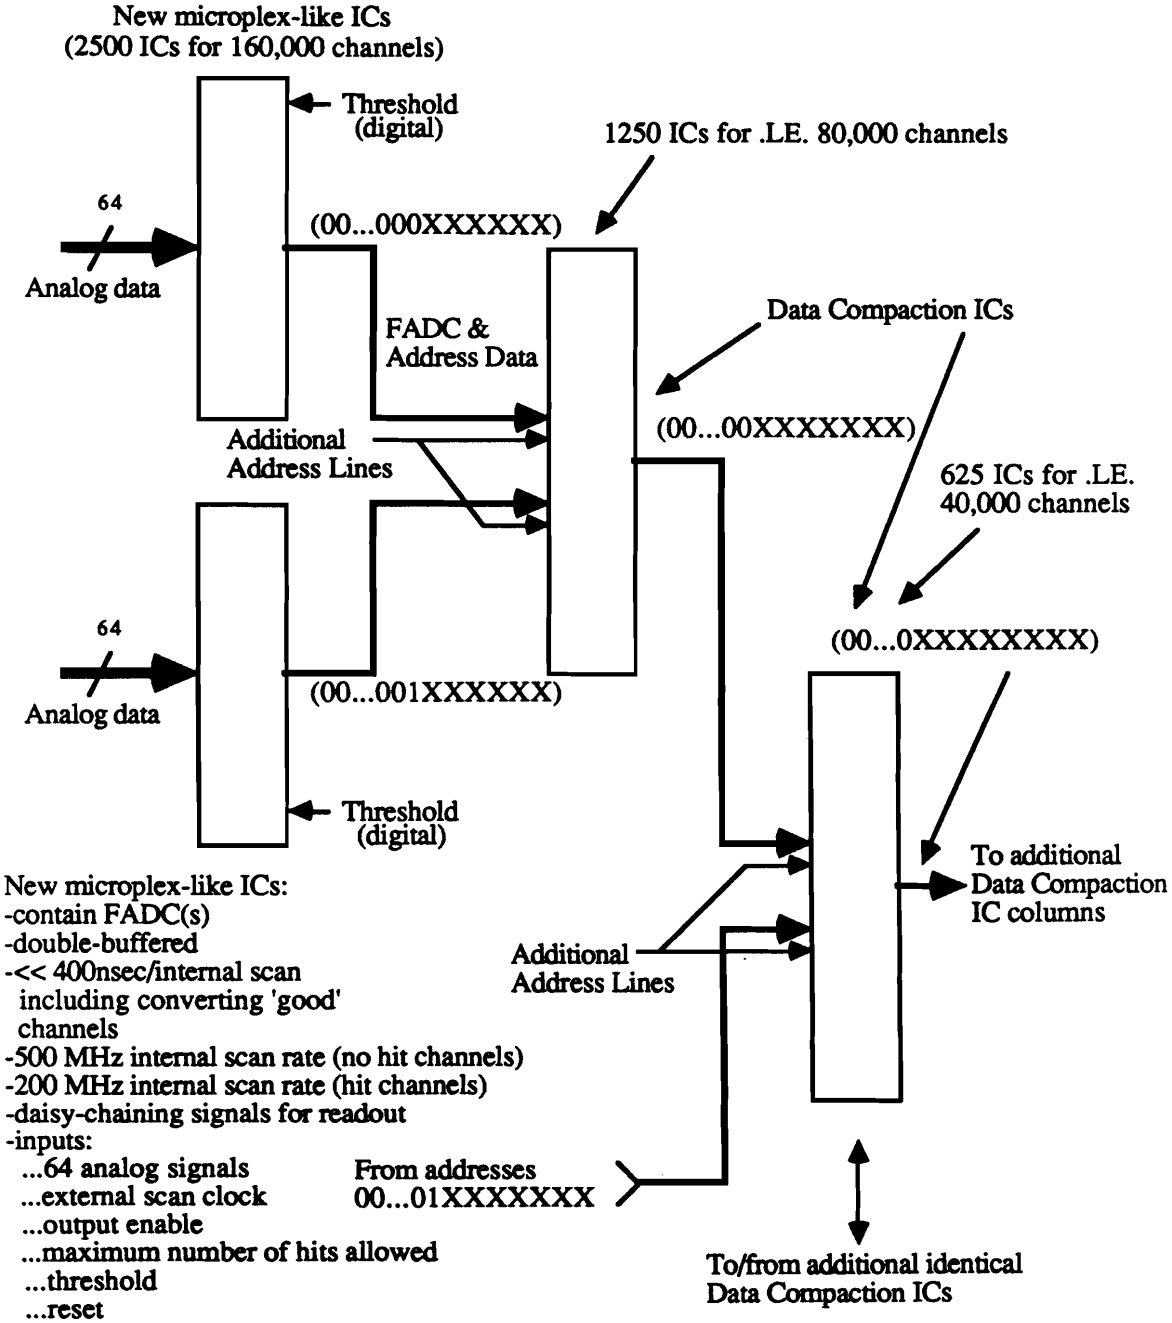

New microplex-like ICs:
-contain FADC(s)
-double-buffered
-<< 400nsec/internal scan
including converting 'good'
channels
-500 MHz internal scan rate (no hit channels)
-200 MHz internal scan rate (hit channels)
-daisy-chaining signals for readout
-inputs:
...64 analog signals
...external scan clock
...output enable
...maximum number of hits allowed
...threshold
...reset
...etc.
-outputs:
...FADC data
...channel address (5 bits)
...overflow (i.e., too many hits)
...end of scan signal
...etc.

Figure 2

B Physics...Front-End Trigger ICs

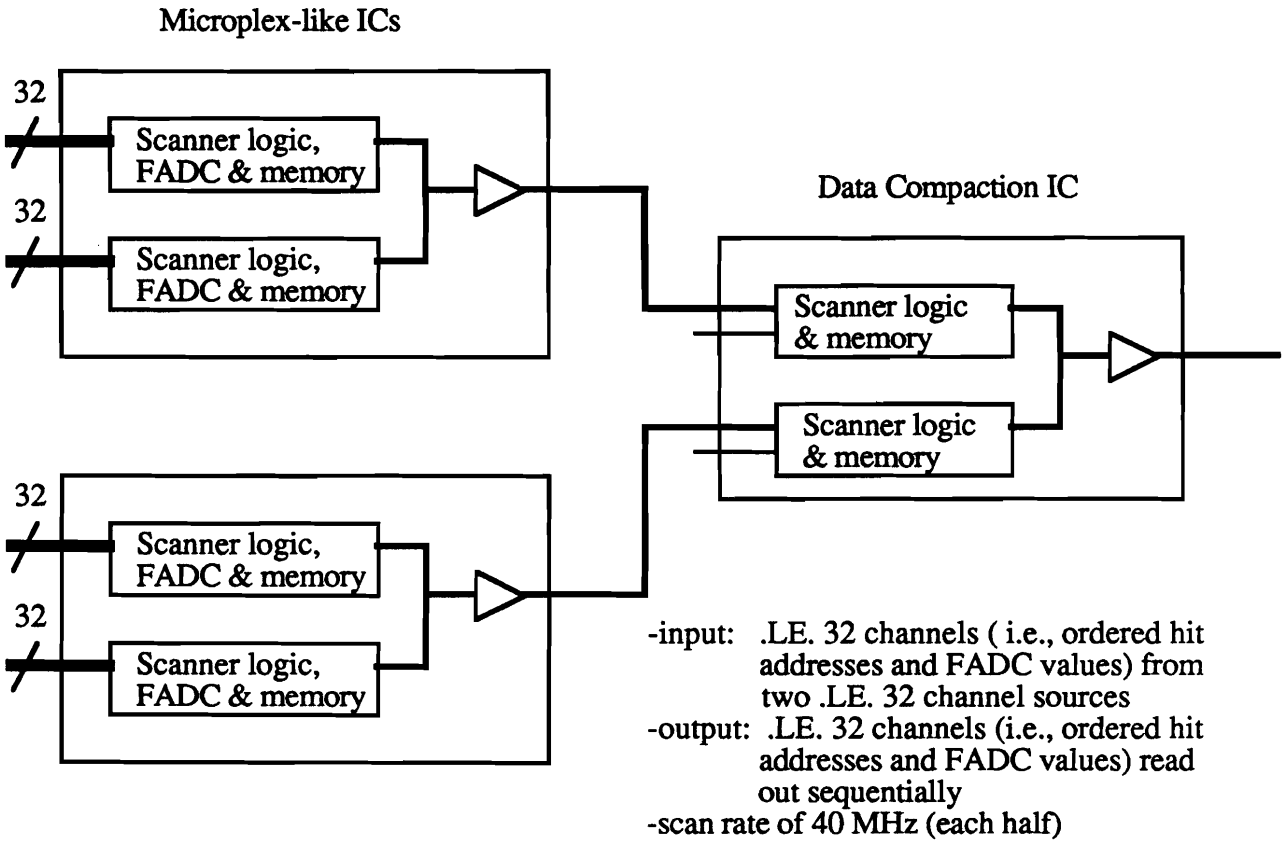

Figure 3

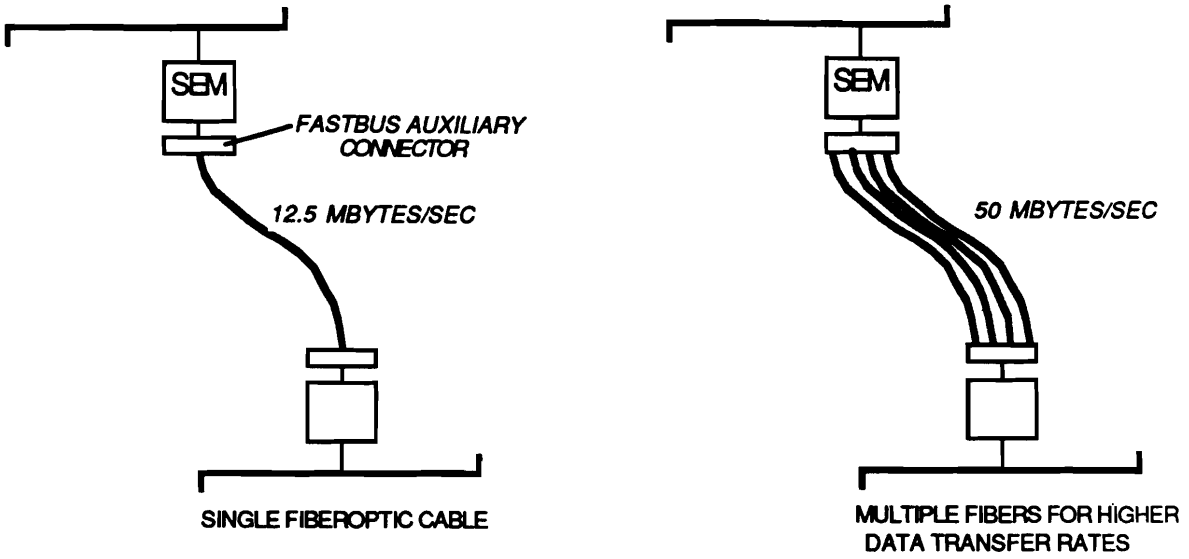

SINGLE SOURCE TO SINGLE RECEIVER

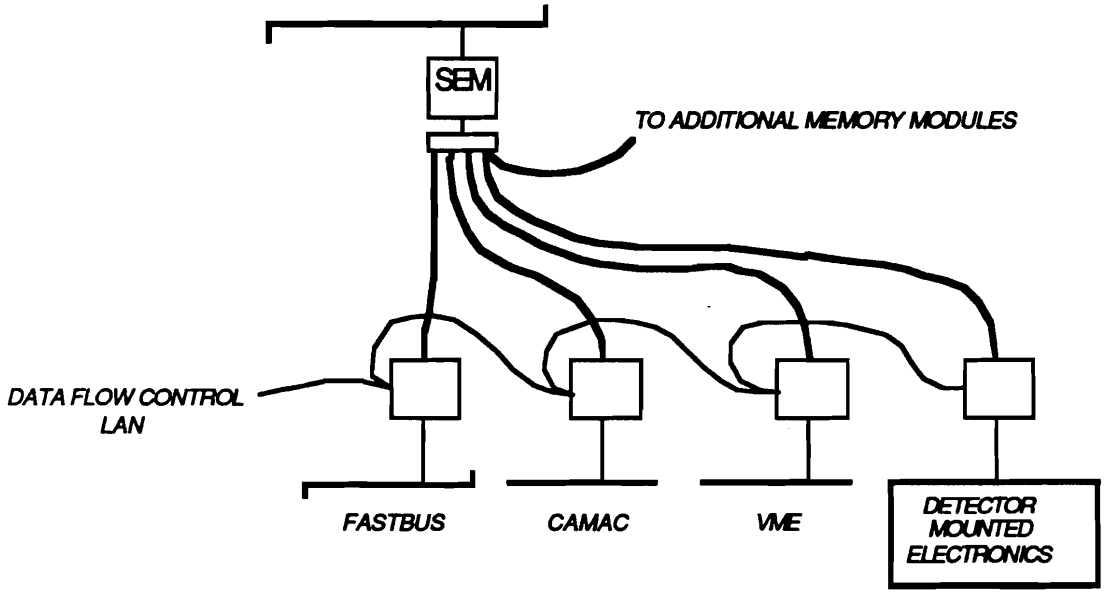

MULTIPLE SOURCES TO A SINGLE RECEIVER

NOTE: ALL IMPLEMENTATIONS USE AMD SUPERNET CHIPS (12.5 MBYTES/SEC) COMMERCIALLY AVAILABLE VLSI FOR DATA TRANSFER APPLICATIONS

FIGURE 4

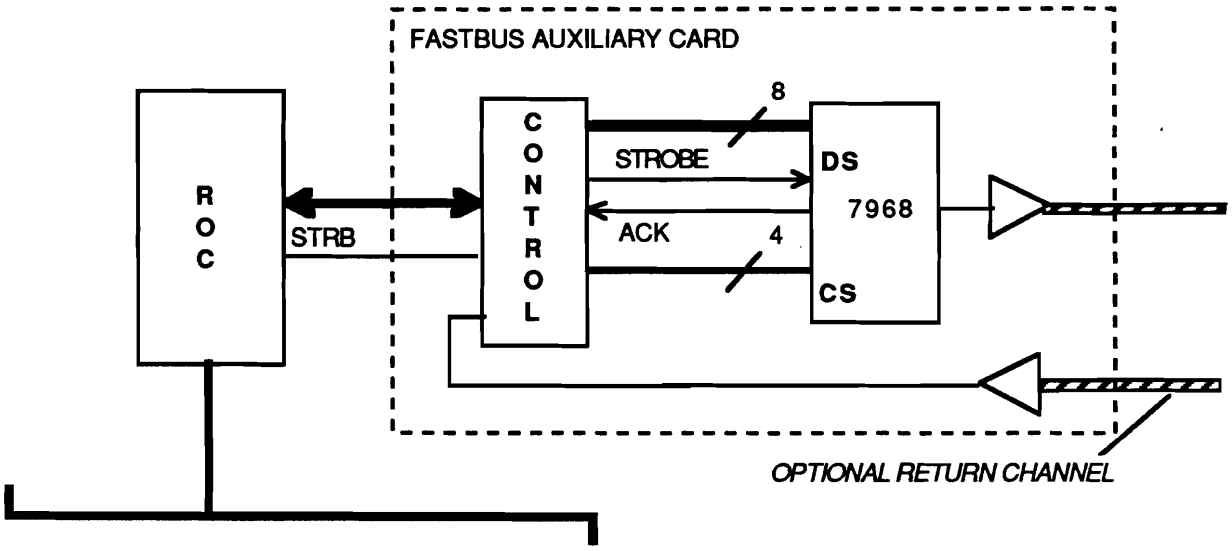

TYPICAL TRANSMITTER IMPLEMENTATION

FIGURE 5

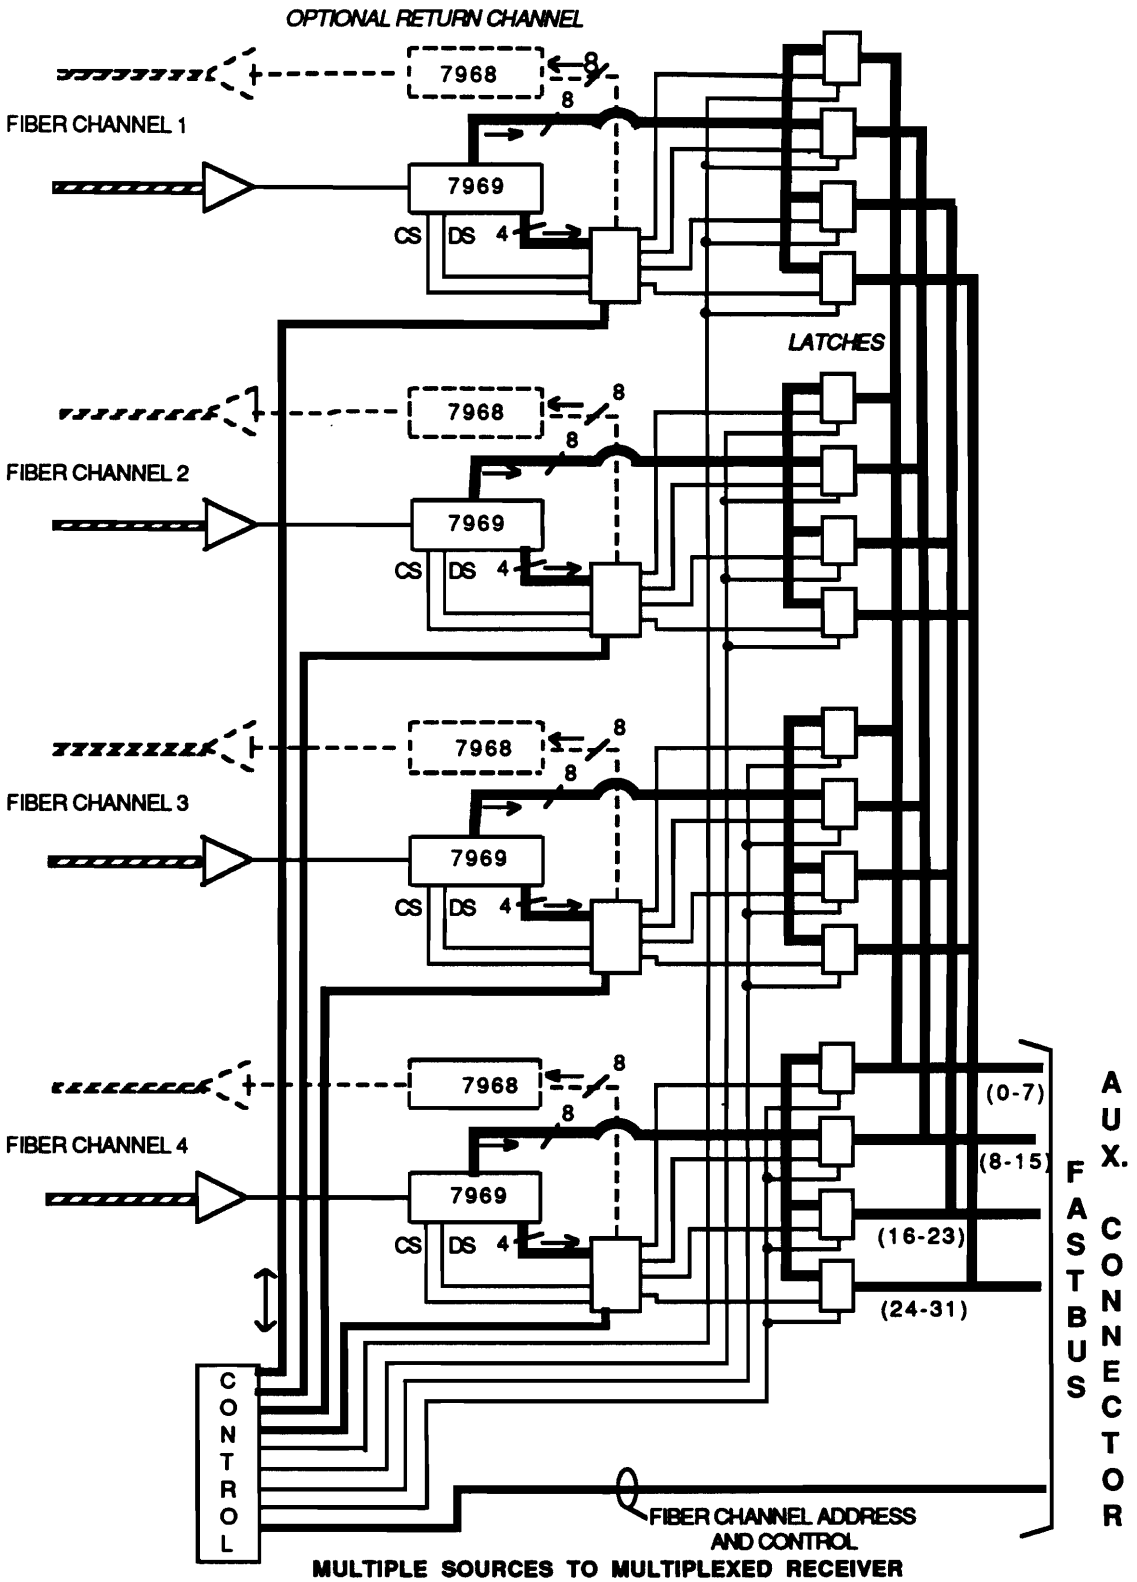

MULTIPLE SOURCES TO MULTIPLEXED RECEIVER

FIGURE 6

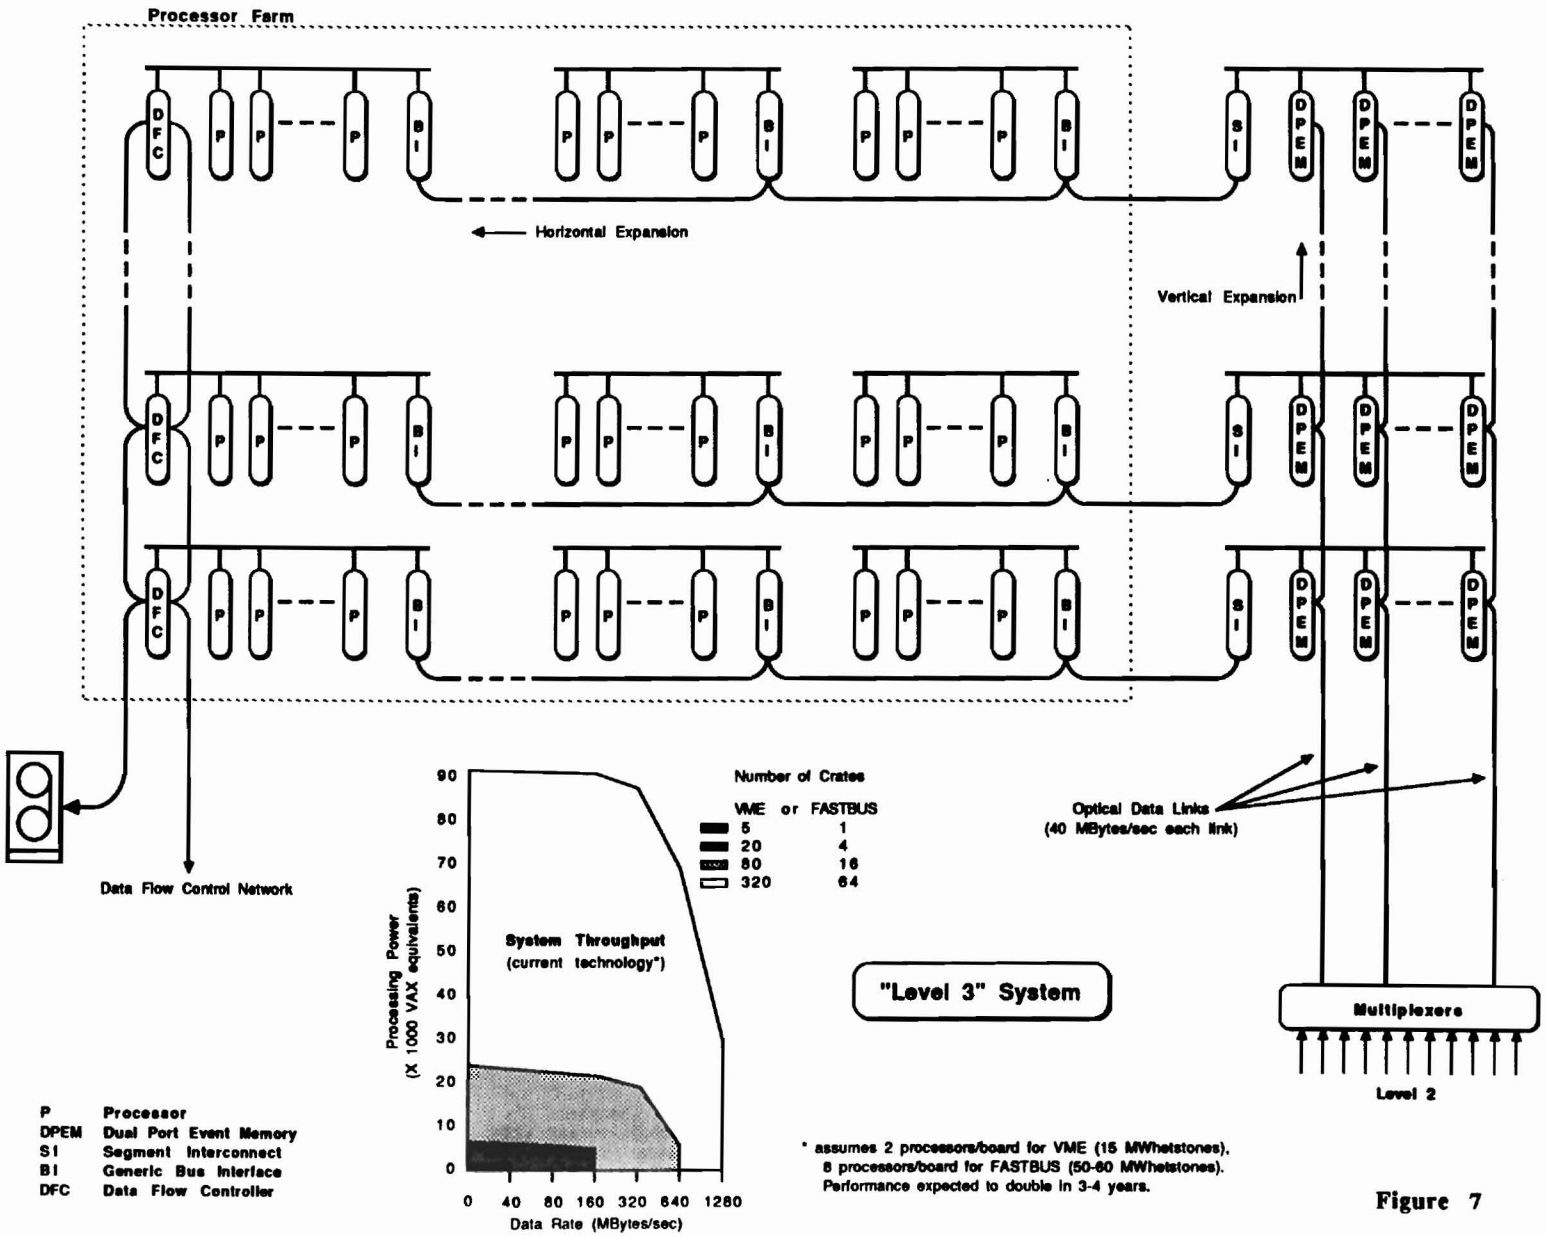
Processor Farm
DFC
P
BI
SI
DPEM
Horizontal Expansion
Vertical Expansion
Data Flow Control Network
Optical Data Links
(40 MBytes/sec each link)
Multiplexers
Level 2
Number of Crates
VME or FASTBUS
5 1
20 4
80 16
320 64
System Throughput
(current technology*)
Processing Power
(X 1000 VAX equivalents)
Data Rate (MBytes/sec)
0 40 80 160 320 640 1280
0 10 20 30 40 50 60 70 80 90
"Level 3" System
P Processor
DPEM Dual Port Event Memory
SI Segment Interconnect
BI Generic Bus Interface
DFC Data Flow Controller
* assumes 2 processors/board for VME (15 MWhetstones),
8 processors/board for FASTBUS (50-60 MWhetstones).
Performance expected to double in 3-4 years.

Figure 7

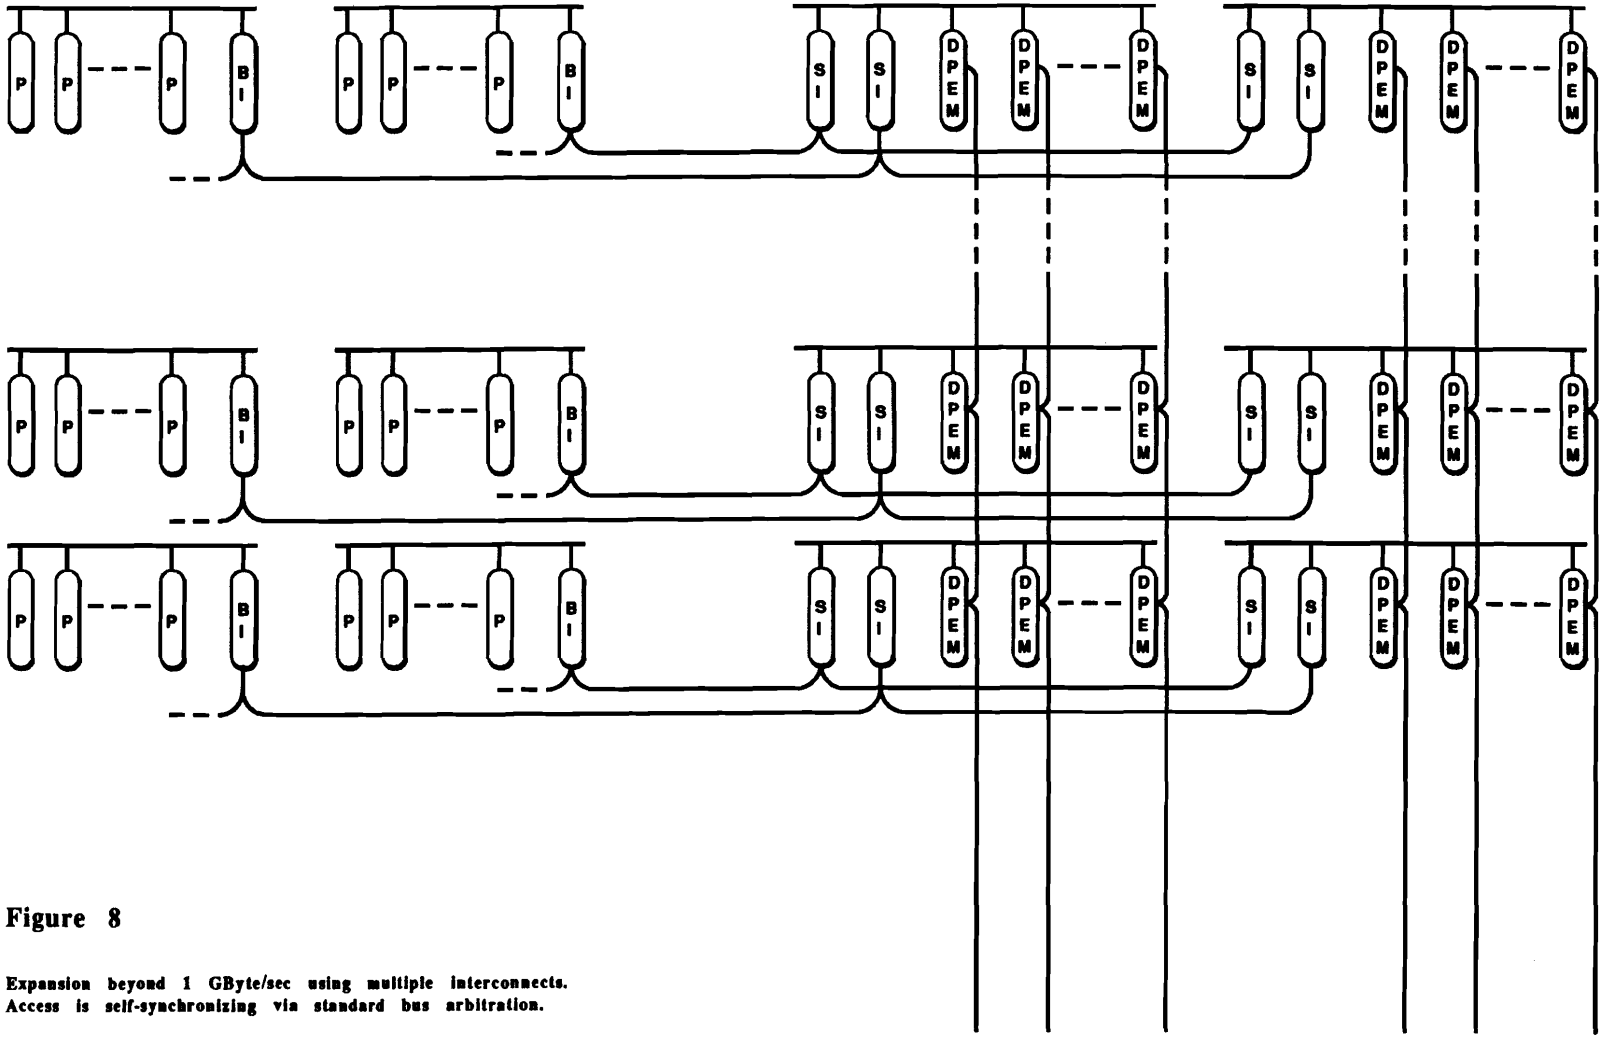

Figure 8

Expansion beyond 1 GByte/sec using multiple interconnects. Access is self-synchronizing via standard bus arbitration.

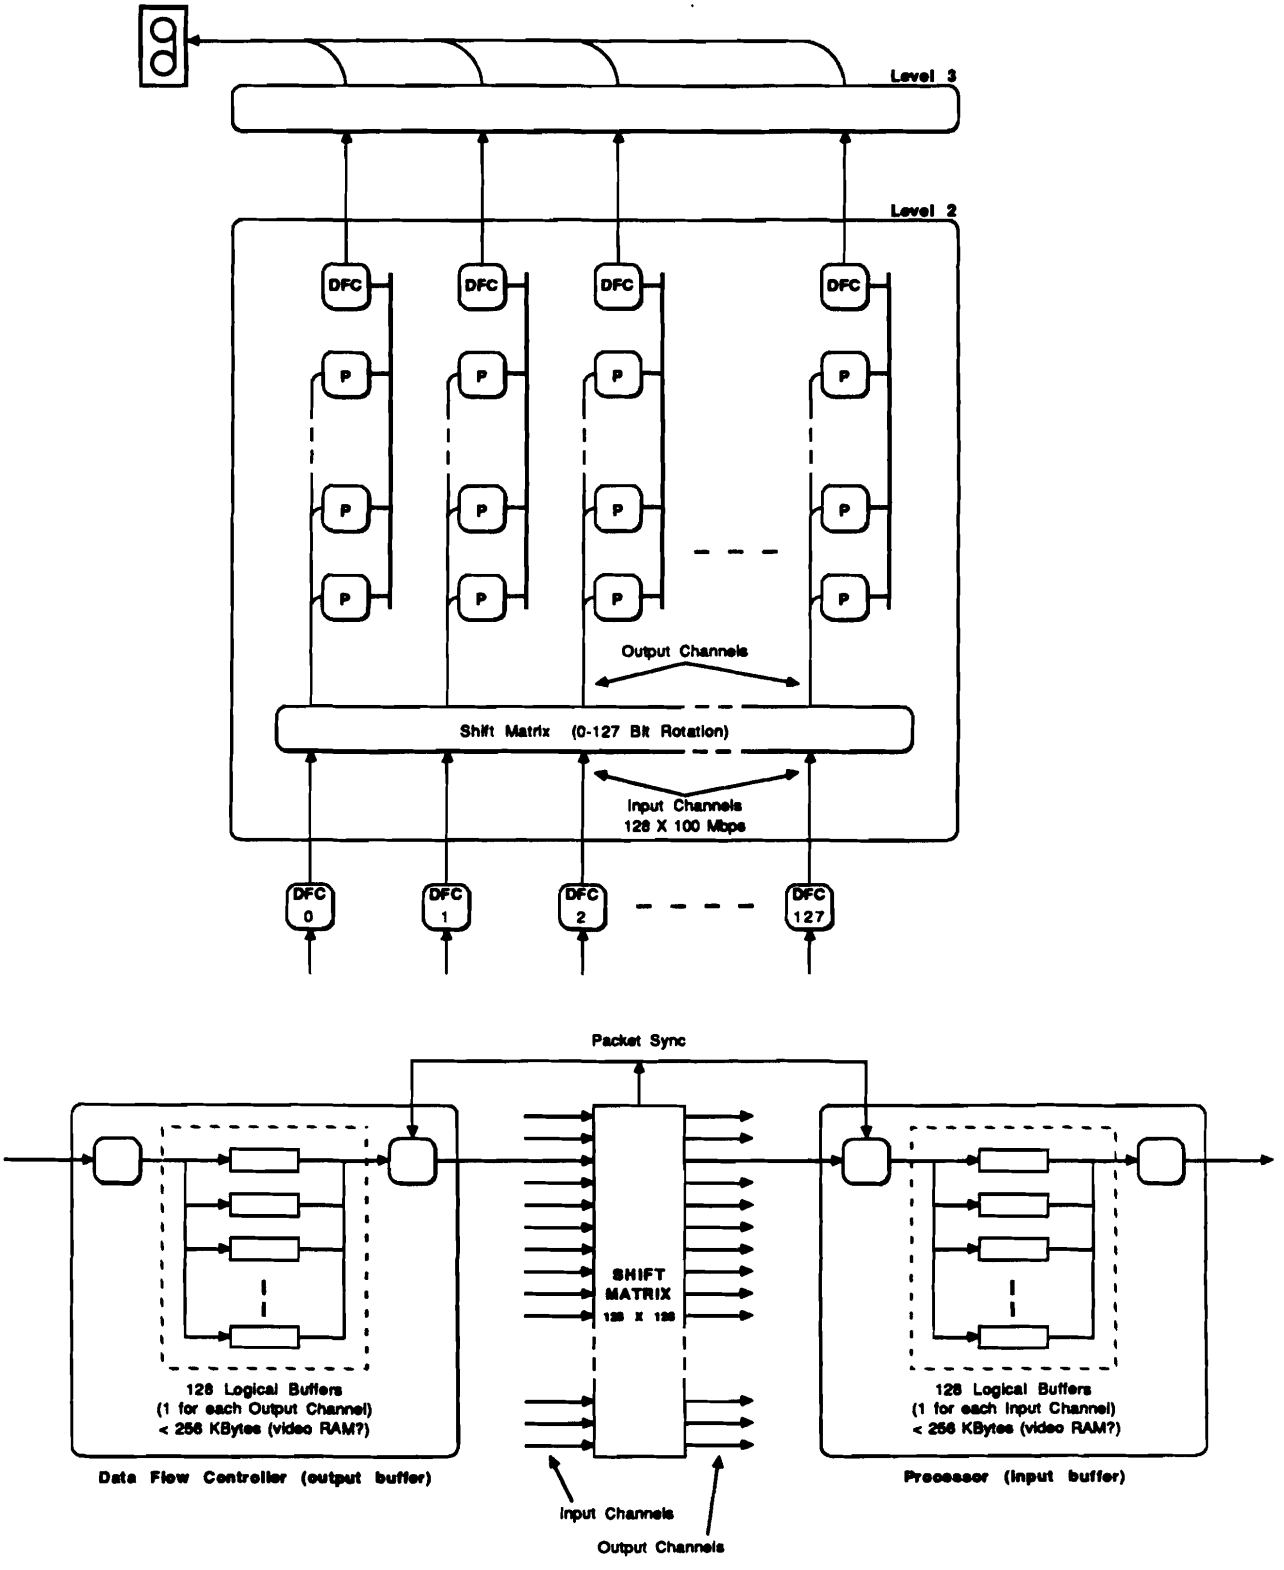

Figure 9 Packet Switching Network (1.6 GBytes/sec)