SDC-92-202

SDC
SOLENOIDAL DETECTOR NOTES

SUPERCONDUCTING SUPER COLLIDER DATA COMPRESSION AND DRIVER/MODULATOR ARCHITECTURE

February 1992

B.J. Cook, K.S. Lackner, D.H. Sharp, C.L. Winter and H.Ziock
Los Alamos National Laboratory

SUPERCONDUCTING SUPER COLLIDER DATA COMPRESSION AND DRIVER/MODULATOR ARCHITECTURE

by
B. J. Cooke, K. S. Lackner, D. H. Sharp, C. L. Winter, and H. Ziock

Los Alamos National Lab
MS-J580
P.O. Box 1663
Los Alamos, N.M. 87545
LA-UR-91-3451

Abstract

We present a prototype architecture that provides chip-level data compression and high-speed optical drive capabilities for rad-hard Superconducting Super Collider data acquisition applications. The data compression is accomplished through a pipelined edge-detection and clustering algorithm, while the driver/modulator uses high-throughput data-management techniques, including variable register lengths, bus queuing, and time-division-multiplexing for fast optical-data transmission. An architecture of this sort can provide efficient transmission of a very large numbers of events.

I. INTRODUCTION

The generic architecture of an inner-detector system is shown in Figure 1. The basic detector unit considered here is a silicon PIN-diode strip detector consisting of 500 to 1000 strips. A particle traversing a wafer will activate a

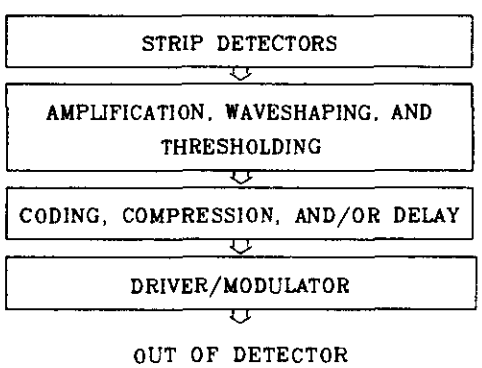

Figure 1. SSC Inner-Detector Architecture.

cluster of strips. The amplified strip signals provide the information on the position of a particle(s) in one dimension. The data acquisition task is to extract the locations of the activated strips and transmit this information to an external data-processing unit that can integrate information of many strip detectors.

Our paper will focus on a pipelined data compression and the driver/modulator architecture of the inner-detector system. Detailed information concerning the application of the data compression and driver/modulator in our prompt/delayed-channel pipelined encoder architecture can be found in [1].

II. DATA COMPRESSION ALGORITHM

The data compression is accomplished by transmitting cluster information instead of information about individual strips. The algorithm proceeds as follows: The raw data are first edge-detected, reducing the clusters to a series of edges. Then, using a pipelined architecture, the relative offsets of the edge boundaries are found. The resulting compressed data vector is then composed of the number of edges, the offsets of each edge from the previous edge, and, if desired, the total number of hits.

An illustration of the algorithm follows. Figure 2 shows a 16 bit (strip) detector with three "clusters" and a binary

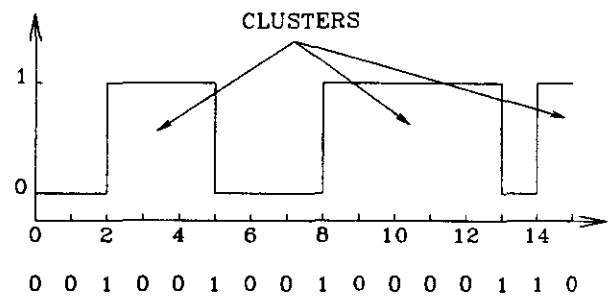

Figure 2. 16-Bit Detector With Three Clusters.

data vector representing the cluster edges. As the data vector is "passed through" the pipelined encoding algorithm, the edges are detected and the edge offsets counted. The decimal encoding of the three clusters in terms of pipeline stage, edge count, and accumulated offset reduce to the following:

Stage	# Edges	Offset
0	0	0
1	0	1
2	1	2/0
3	1	2/1
4	1	2/2
5	2	2/3/0
6	2	3/2/1
7	2	2/3/2
8	3	2/3/3/0
9	3	2/3/3/1
10	3	2/3/3/2
11	3	2/3/3/3
12	3	2/3/3/4
13	4	2/3/3/5/0
14	5	2/3/3/5/1/0
15	5	2/3/3/5/1/1.

Note that the "/" denotes an edge boundary and that it also needs to be encoded. The equivalent binary representation of the data vector in Figure 2 is

Stage	Accumulator	# Edges	Data Vector
0	[Offset] 0 [Edge] 0	0	0010010010000110
1	1 0	0	010010010000110
2	10 01	1	10010010000110
3	101 010	1	0010010000110
4	1010 0100	1	010010000110
5	1011 0101	10	10010000110
6	10111 01010	10	0010000110
7	101110 010100	10	010000110
8	101111 010101	11	10000110
9	1011111 0101010	11	0000110
10	10111110 01010100	11	000110
11	10111111 01010100	11	00110
12	101111100 010101000	11	0110
13	101111101 010101001	100	110
14	1011111011 0101010011	101	10
15	10111110111 01010100110,	101	0

where the input data vector is in the right hand column and the accumulated edge boundary and offset are on the left. The edge location and offset are now represented by two separate binary vectors. Note that a single vector element can be processed per pipeline stage allowing the algorithm to be implemented with a 60-MHz system clock (we are hard-pressed, as it is, to find a low-power, rad-hard CMOS process capable of 60-MHz operation, let alone a higher clock rate). Also, note that the compressed data is right (or left) justified.

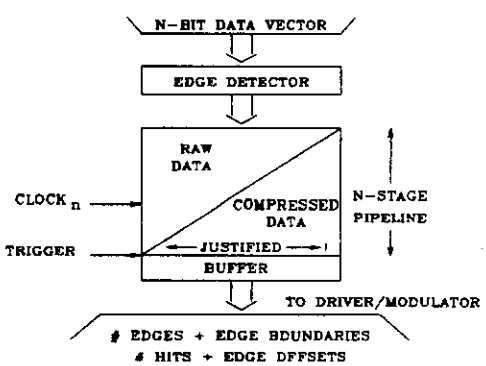

Figure 3. Compressor Architecture.

III. COMPRESSOR AND DRIVER/ MODULATOR ARCHITECTURE

A prototype compressor architecture is shown in Figure 3. The compressor would accept an N-bit data vector, edge detect, and it compresses the edge data utilizing the previously described compression algorithm. The edge detector and compressor can be implemented with relatively straight-forward digital logic and shift registers for the required pipelined architecture. The resulting two output vectors would then be composed of the number of edges + edge boundaries and the number of hits + edge offsets.

Typically, since N is on the order of 500-1000 bits and realistic chip Input/Outputs (I/O) are on the order of 100-200, input vector segmenting is required. A practical realization of a four-segment 4x128 (512-bit) compressor and driver/modulator architecture is shown in Figure 4. The chip

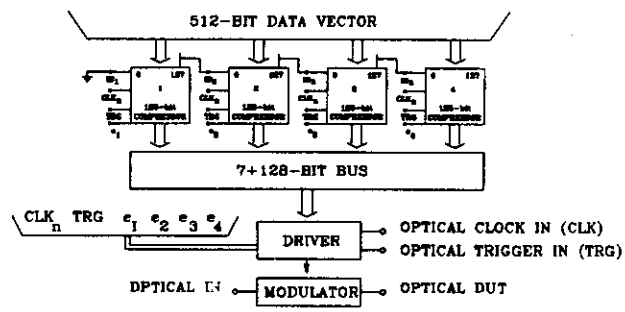

Figure 4. 4x128-Bit Compressor and Driver/Modulator.

set is composed of four 128-bit compressors. On arrival of the external trigger, the compressed data is multiplexed over a 7+128-bit data bus to the driver/modulator chip for optical transmission. While a segmented architecture allows the implementation of chips with reasonable I/O pin counts, data compression is degraded as will be discussed in Section IV. Figure 5 shows the corresponding driver/modulator architecture for the 4x128 compressor. The 4x(7+128)-bit vector segments are sequentially read and stored in the bus buffer registers. After the bus buffer acquires all four data segments, a fast block transfer of the segments is made to the time-division-multiplexing (TDM) buffer where each segment is then sequentially shifted into the TDM variable length drive register for optical transmission. The TDM variable length drive register provides a means of sending only non-zero compressed information (determined by the edge-count vector), preventing the transmission of redundant information. The 2:1 TDM modulator takes advantage of the equal lengths of edge boundary and offset vectors. Note that an optical clock and trigger are used to minimize time jitter.

IV. MONTE CARLO SIMULATION

The compression efficiency of the algorithm for single and multiple segmented compressors was evaluated with a simple a Monte Carlo simulation. The simulation observed the following steps:

(i) Random selection of the positions of clusters over the width of the detector wafer.

(ii) Random selection of the width of each cluster.

(iii) Calculate the encoded binary representation.

(iv) Repeat i-iii to determine the average compression ratio.

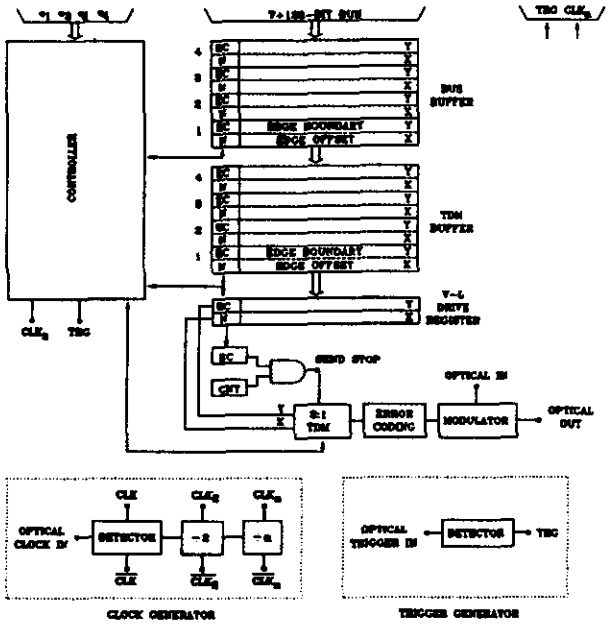

Figure 5. 128-Bit Driver/Modulator Architecture.

A uniform distribution is assumed for (i) and (ii). In addition, the clusters are composed of 1 to 7 hit strips and are allowed to overlap. The average compression ratio was defined as

$$\frac{\text{Total Number of Strips}}{\langle \text{Offset Vector} + \text{Boundary Vector} \rangle},$$

where <> indicates averaging over 25 samples. Figures 6 and 7 show the compression ratios for a single-segment and a 4-segment 512-bit compressor. The single-segment compressor yields a 15:1 compression for a single cluster and 6:1 for five clusters, while the 4-segment provides a 7.5:1 and 4.5:1 compression ratio for one and five clusters, respectively. The degraded efficiency is caused by the reduction in binary encoding efficiency with decreasing vector size and encoding clusters which overlap segment boundaries (such clusters require two or more binary numbers to encode). Hence, compression efficiency is penalized as the number of segments are increased and data vector lengths are reduced. Techniques to improve segmented compression by combining segments before transmission are under investigation. Note that the sum of the offsets in a segment will always add up to the segment size, thus providing an end-of-transmission sum-check.

V. CIRCUIT DESIGN CRITERIA

This section considers various design criteria associated with the practical implementation of the compressor and driver/modulator architecture.

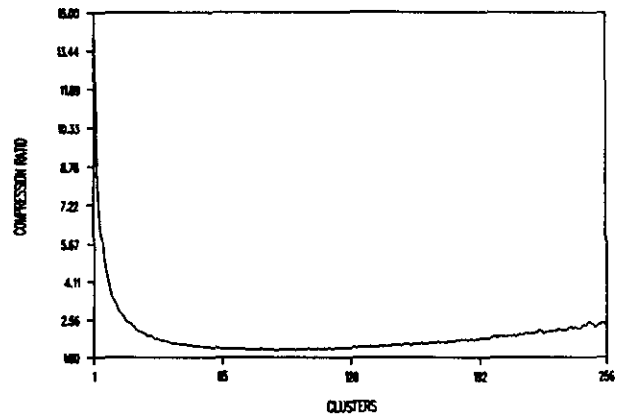

Figure 6. Compression Ratio for a Single-Segment, 512-Bit Vector.

A. Compressor Register Requirements

The minimum number of registers required to implement a 512-bit compressor architecture for 1 to 8 segments assuming a 60 MHz clock rate, 1-shift/cycle, and 1-add/cycle is as follows:

Segment Size	# Registers	Compression
512	393K	15:1
256	196K	10.5:1
128	98K	7.5:1
64	49K	5:1.

Note that the pipeline delay, compression efficiency, bandwidth (the product of segment size times 60-MHz clock rate), and number of registers (i.e. circuit density and power dissipation) can be traded off. Again, emphasis should be placed on the fact that segment size will be limited, by practical engineering considerations, to I/O pins counts of 100-200.

An estimate of the approximate power dissipation per register can now be obtained. Using the 4x128 compressor chip discussed in section III, there are 25K registers required per chip. If a total of 250 mW of power is allocated for the data compression of 512 bits, then each chip is limited to approximately 63 mW, or leaving less than 2.5μW per register not including overhead chip circuitry. Note that the compressor architecture requires 50% more registers than does a simple delay line architecture.

B. Driver and Optical Modulator

The driver/modulator chip can be implemented with various technologies depending on cost, speed, power dissipation, and complexity. For our SSC applications, we are considering three combinations of driver-modulator topologies: Silicon-$LiNbO_3$, GaAs-$LiNbO_3$, and GaAs-GaAs. The silicon and GaAs drivers devices would be bipolar and MESFET respectively while the $LiNbO_3$ and GaAs modulators are both electrooptic waveguide devices [2].

A summary of the three driver/modulator technologies under consideration is as follows:

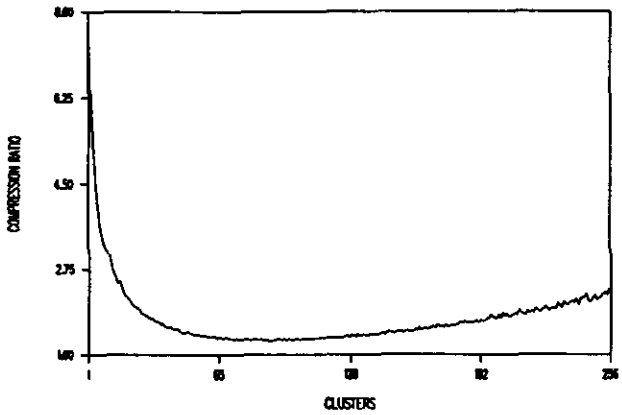

Figure 7. Compression Ratio for a Four-Segment, 512-Bit Vector.

Driver-Modulator	Speed	Technology
Silicon-$LiNbO_3$	Moderate	Mature
GaAs-$LiNbO_3$	Moderate+	New
GaAs-GaAs	Fast	Under Development.

The first two technologies are hybrid devices where the $LiNbO_3$ modulator is wire bonded to the silicon or GaAs substrate while the third integrates the driver and modulator on the same substrate, reducing parasitic losses and potentially reducing production costs. The best power-speed figure is realized with GaAs High-Electron-Mobility-Transistors (HEMT) [3]. However, it should be noted that HEMT processes are still primarily confined to laboratories and not readily commercially available. In addition, GaAs chips in general exhibit a lower thermal conductivity then its silicon counterpart.

Our preliminary design goal is the development of a silicon-$LiNbO_3$ driver/modulator operating at 1-5 GHz with a power dissipation of 250-500mW (if this goal is unattainable we will then focus on an HEMT driver). The optical modulator prototype is a Mach-Zehnder interferometer [2] adapted for a low drive digital voltage [4]. A LANL radiation exposure test of the $LiNbO_3$ Mach-Zehnder interferometer (total dose 10^{14} 647 MeV protons/cm^2) [5] along with the appropriate selection of optical fibers [6] indicate that good resistance to SSC level radiation exposure can be expected.

VI. CONCLUSION

We have presented a prototype architecture that provides chip-level data compression and high-speed optical drive capabilities for rad-hard Superconducting Super Collider data acquisition applications. We have discussed how data compression is accomplished through clustering and introduced a flexible high-speed driver/modulator for fast optical-data transmission. A practical 4x128 compressor architecture was introduced to illustrate the algorithm. We believe an architecture of this sort can provide efficient transmission of very large numbers of events.

VII. REFERENCES

[1] B. J. Cooke, K. S. Lackner, D. H. Sharp, and C. L. Winter, "Efficient Data Transmission From Silicon Wafer Strip Detectors," Nuclear Science Symposium, Santa Fe, NM, Nov. 5-9, 1991.

[2] R. C. Alferness, "Waveguide Electrooptic Modulators," IEEE Transactions on Microwave Theory and Techniques, Vol. MTT-30, No. 8, August 1982.

[3] M. Abe, T. Mimura, K. Nishiuchi, A. Shibatomi, M. Kobayashi, and T. Misugi, "Ultra-High-Speed HEMT Integrated Circuits," Semiconductors and Semimetals, Vol. 24, Academic Press, 1987, pp. 253-254.

[4] J. L. Jackel, and J. J. Johnson, "Nonsymmetric Mach-Zehnder Interferometers Used as Low-Drive-Voltage Modulators," Journal of Lightwave Technology, Vol. 6, No. 8, August 1988, pp. 1348-1351.

[5] Private Communication H. Ziock Los Alamos National Lab and M. Lowry Lawrence Livermore National Lab.

[6] B. D. Evans and G. H. Sigel, Jr., "Radiation Resistant Fiber Optic Materials and Waveguides," IEEE Transactions on Nuclear Science, Vol. NS-22, No. 6, December 1975, pp. 2462-2474.