[See page 425 for Figure 3.]

SESSION XX: SPECIAL PURPOSE ACCELERATORS

FAM 20.4: A Dynamically Reconfigurable Interconnection Chip

Chi-Yuan Chin, Wen-Tai Lin, Juh-Ping Hwang, Sow Chu, Glenn Forman, Robert Dunki-Jacobs, Steven Karr, John Mallick**, H.T. Kung***, Allen Sussman***, F.H. Hsu***, T. Nishizawa***

General Electric Corporate Research and Development Center

Schenectady, NY

DYNAMIC RECONFIGURATION OF INTERCONNECTION is important in today's multiprocessing computing systems. A reconfigurable network between processing elements is essential for systolic array, vector processing, or pipeline processing systems. This paper focuses on the development of a Link and INterconnection Chip (LINC), an interface chip optimized for dynamically reconfiguring the data flow and sequencing the processing elements (multipliers, memory, ALU) of a computing system. With the LINC chip, one can optimize the data flow and sequencing required to compute each algorithm executed in a system application. Key functions include data synchronization, computational resource scheduling and compensation for variable pipeline latency. In addition, the control memory must permit the down load of new execution sequences concurrent with normal operation.

The LINC is shown in Figure 1. A 8x8, 4b wide crossbar network serves as an efficient link between system functional modules, such as arithmetic units, register files, and I/O ports. The FIFO, Programmable Delay Register (PDR), and Pipeline Register File (PRF) adjust data stream delay cycles to facilitate the synchronization scheduling task for both data and computational resources. Using pattern commands, which are prestored in the RAM, the chip can configure the interconnection of its 32 inputs, 32 outputs and delays for many applications, such as, optimization, signal processing, sorting, data shuffling (e.g., corner turning), without interrupting data flow. More over, the LINC can be used in a bit-sliced fashion to form interconnections with datapaths wider than 4b. By tri-stating the proper data output pins, multiple copies of the chip can form crossbar interconnections larger than 8x8.

The chip, as shown in Figures 1 and 2 has eight 4b datapaths consisting of the 8x8 crossbar interconnection, either a FIFO or a programmable delay for each of its inputs, and a pipeline register file for each of its outputs. The connection pattern between sources and destinations (including broadcasting) is determined by the upper 24b in a Command Pattern Register (CPR) which is 64b wide. At each clock cycle, the CPR is updated from one of the two memory banks, which has been loaded from the host. Each bank can be pre-stored with 32 different commands. The advantage of two banks is to allow users to update one bank of control patterns without interrupting data flow through the chip.

The FIFO/PDR consists of a 4b wide, 31-stage shift register, Each shift register may be operated as either a Programmable Delay Register (PDR) or part of a FIFO, as determined by a dedicated pre-loaded register (8b wide Delay Code Register). Input data to the chip will be loaded in the PDR at the stage determined by the lower 5b of the delay code register, while the data already loaded in the PDR will be shifted to its next stage at each clock cycle. The data at the first stage will be shifted to the crossbar network. The delay needed for data synchronization is accomplished in this fashion. Alternatively, the 31-stage input register can be used as a FIFO. In this mode each FIFO (A or B) is controlled interactively through its dedicated read and write request signals. The status (FIFO full or empty) of each FIFO is reported at each cycle. Since each FIFO/PDR can be assigned as part of FIFO A or FIFO B, the FIFO width can be preset from 0b to 32b, depending on the requirement of the computing system.

The Pipeline Register File (PDR) consists of a 4b wide, 14-stage shift register which has its inputs connected to the crossbar. Under program control the user can non-destructively read out the data from any of the 14 stages. The output data of the crossbar can be written to the external pins or loaded into the Pipeline Register File for later readout, as determined by the remaining lower 40b stored in the CPR. In general, the FIFO/PDR is used to delay signals to provide a synchronizing function while the Pipeline Register File is used for any data which must be reused in the system because of its non-destructive readout capability.

The data path in LINC consists of two primary pipeline segments: one includes input pads and FIFO/PDRs, and the other includes a crossbar and output pads. An on-chip 4K static RAM designed using a six-transistor cell is addressed either sequentially through an 5b address counter or randomly through control pins. The design concept for each output channel of the crossbar network is similar to an 8-to-1 selector, decoded by a set of 3 control bits. Thus, the broadcasting capability is built in, and destination conflict is avoided.

The chip has been fabricated in 1.25μm double-metal CMOS Bulk technology. It contains over 70K transistors occupying 6.2mm by 7.6mm as shown in Figure 3. The device, packaged in a 104-pin grid array, has 32 inputs, 32 outputs, and 30 control signal and status bits. To facilitate test and debug, scan paths have been included at the input and output of the crossbar,

---

***Current address: Philips, Eindhoven, Netherlands.**

****Current address: Union College, Schenectady, NY.**

*****Current address: Department of Computer Science, Carnegie-Mellon University, Pittsburgh, PA.**

[1] **Lin, W.T., Chin, C.Y., "A Reconfigurable Processor Array Using LINC Chip",** *Proceedings of the International Systolic Workshop, Oxford University*, **England; July, 1986.**

[2] **Hsu, F.H., Kung, H.T., Sussman, A., and Nishizawa, T., "Architecture of Link and Interconnection Chip",** *Proceedings of Chapelhill Conference on VLSI*, **p. 439-462; 1985.**

0193-6530/87/0000-0276$01.00 © 1987 IEEE

between the RAM and the control pattern register, and at the eight delay code registers. In this way the design of each module was verified independently. The control signals for all data paths also can be observed. When operated with the input stage configured as a programmable delay register, a data transfer rate of 25MHz has been measured at room temperature. While a FIFO is assigned, the rate drops to 15MHz because of delay from the input pads and the ripple counter.

Operating at a cycle time of 40ns, this chip makes it possible to interconnect a variety of high-performance processing elements with reduced board complexity. This reduction of chip count is especially significant for implementations of multiprocessors such as systolic arrays, as shown in Figure 4, which call for large numbers of processing elements[1]. In the graphics geometry system, 32 LINC chips along with several 25MHz processing elements were configured to achieve a throughput of approximately 22.5K vertices in one frame period (30ms)[2]. Furthermore, other applications have been developed, e.g., computed tomography imaging and simulated annealing CAD engine, which makes use of dynamic reconfiguration the LINC chip provides.

Acknowledgments

The authors would like to acknowledge the technical support provided by D. Locker, D. Botticello, K. Rohling, and engineers from VLSI Technology Laboratory.

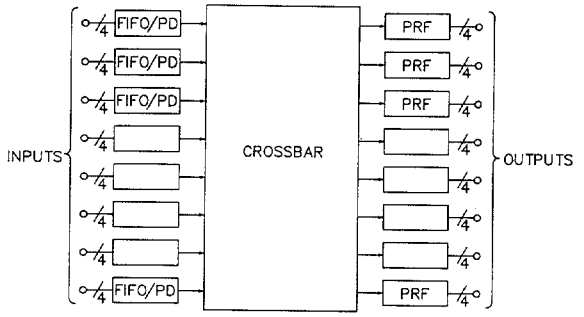

FIGURE 1—Simplified block diagram.

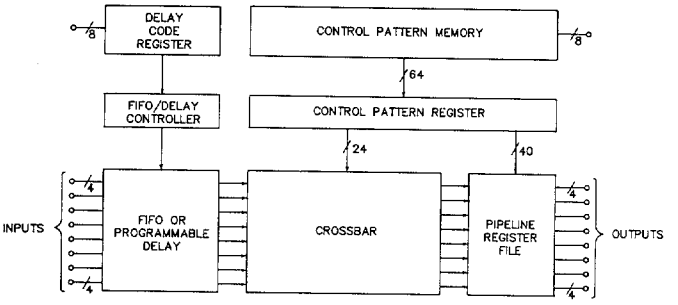

FIGURE 2—Detailed block diagram.

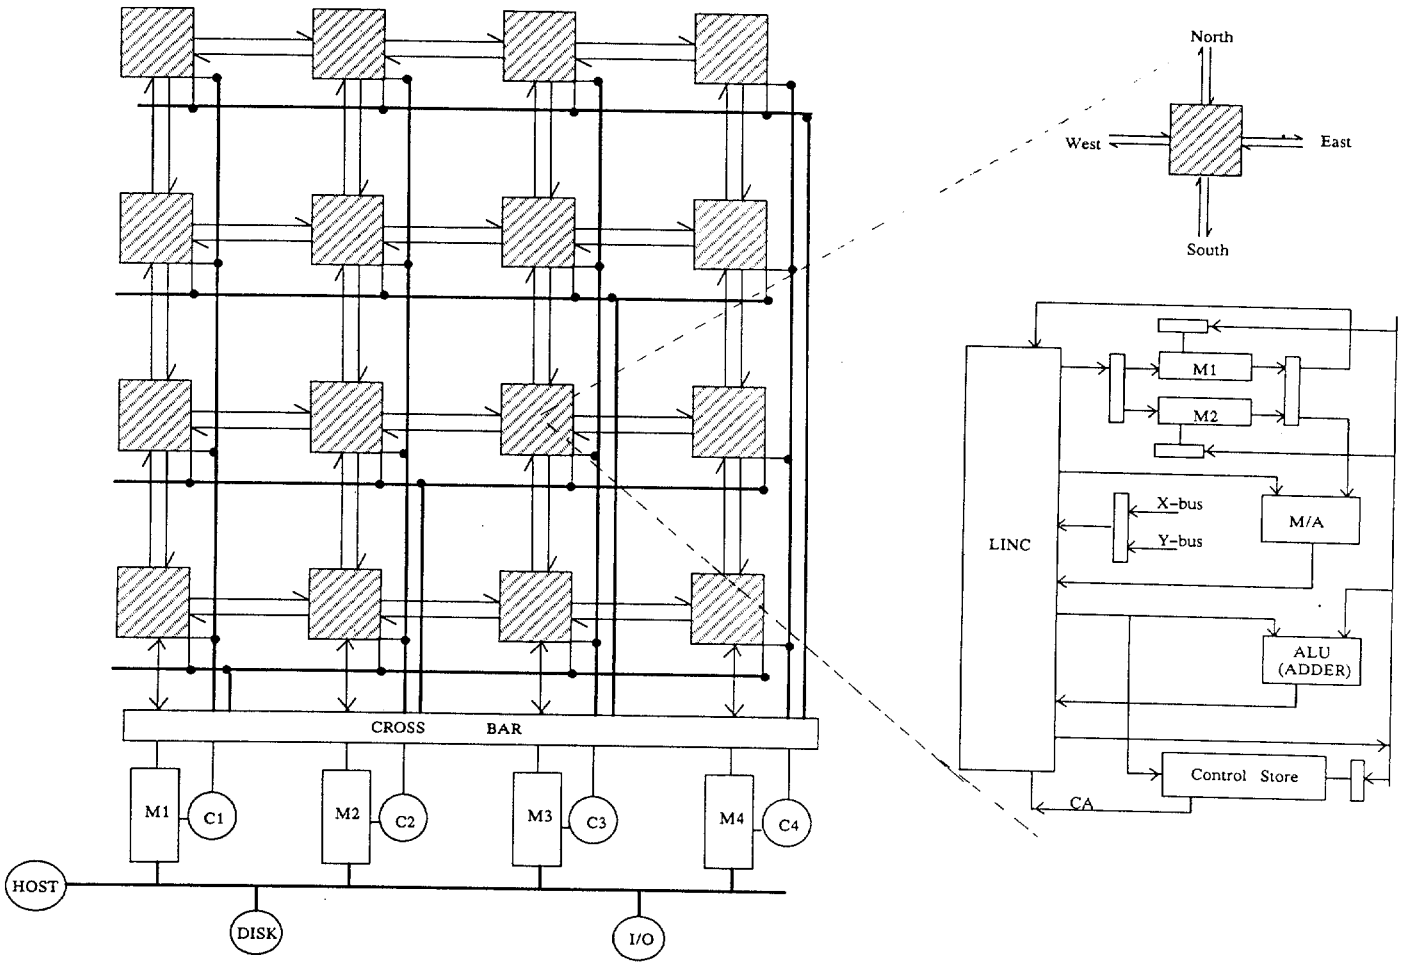

FIGURE 4—System architecture.

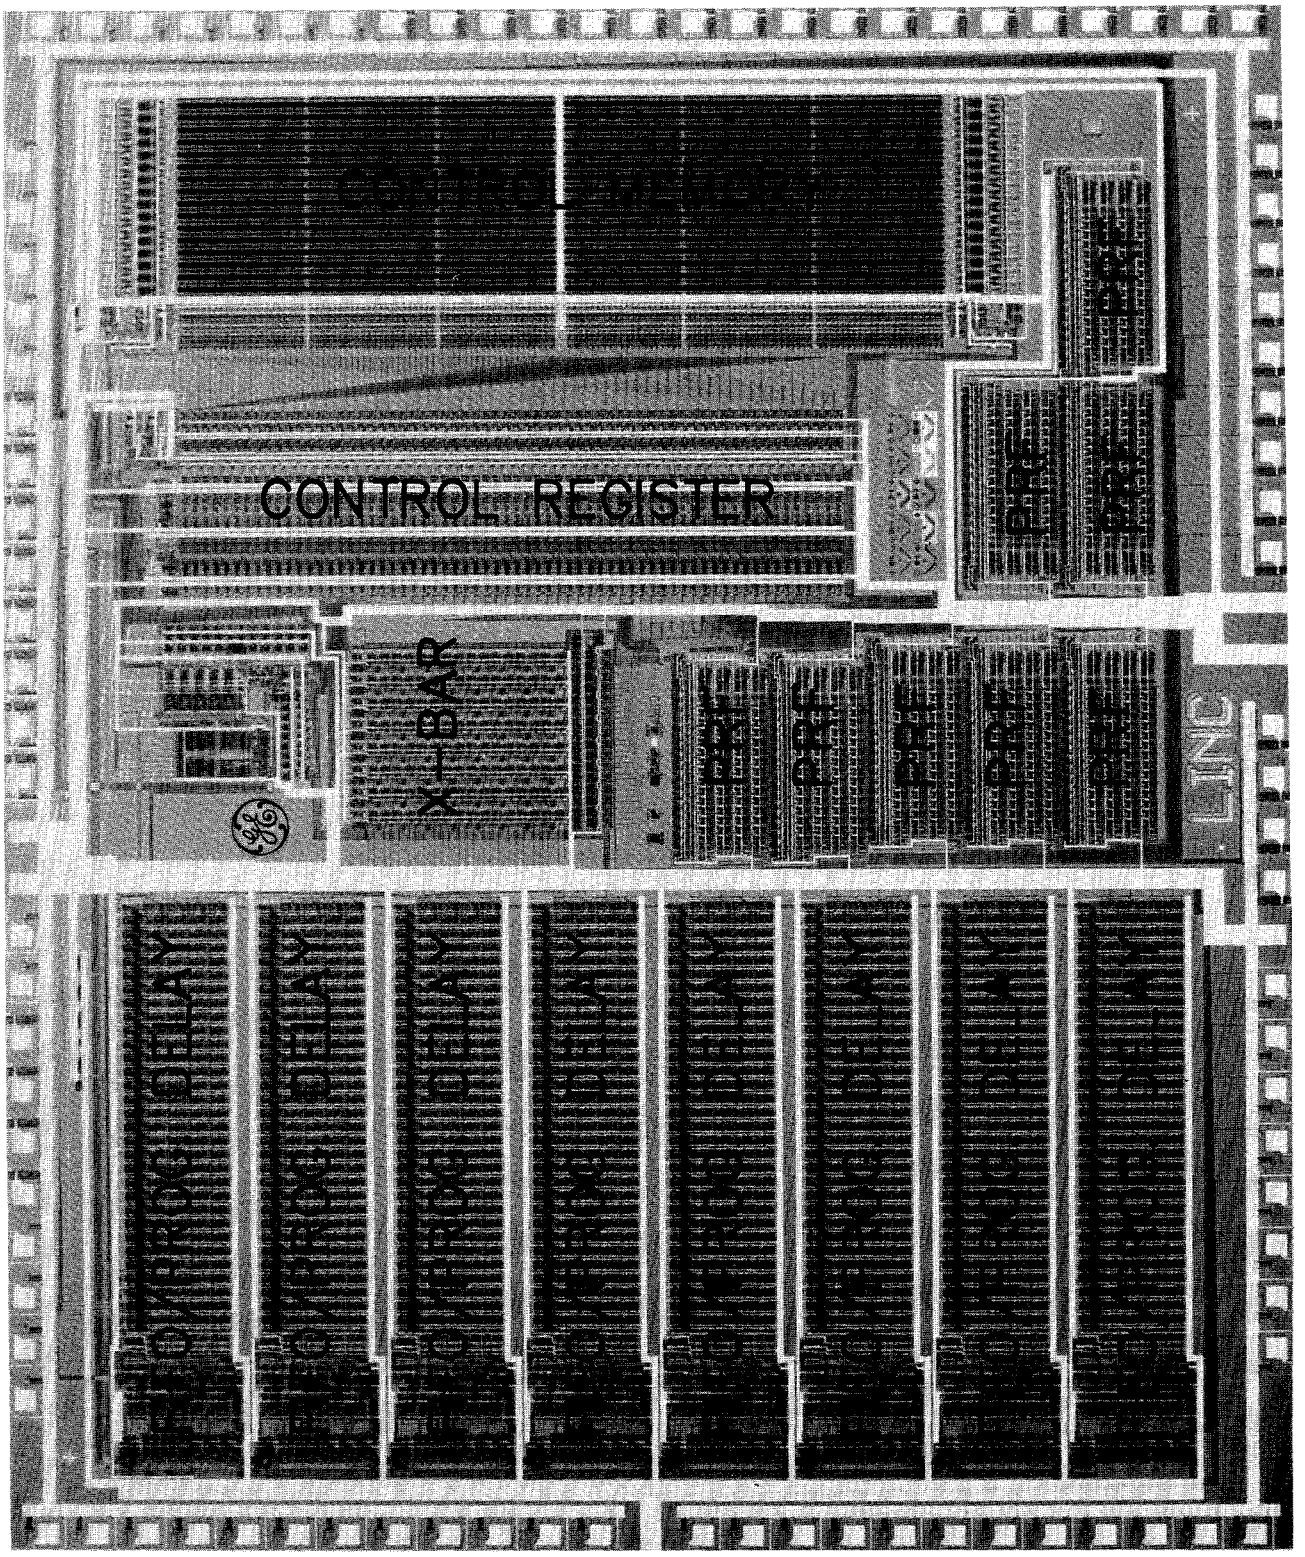

FIGURE 3—Photomicrograph of LINC chip.