

**The ICAP Parallel Processor
Communications Switch**

Deepak Rana, Charles C. Weems

**Computer and Information Science Department
University of Massachusetts**

COINS Technical Report 89-02

**This work was funded in part by the Defense Advanced Projects Agency
under contract number DACA76-86-C-0015, monitored by the U.S. Army
Engineer Topographic Laboratory.**

ABSTRACT

The architecture of a custom VLSI Parallel Communications Switch (PARCOS) chip is described. The PARCOS chip consists of a communication matrix of 32 bit serial inputs and 32 bit serial outputs and an on-chip control memory. The control memory called the Connection Pattern Cache (CPC), is constructed so that PARCOS can hold up to 32 of the most frequently used connection patterns between its inputs and outputs. Any of these stored patterns is incrementally modifiable, and the connection pattern of the communication matrix can be switched from one stored pattern in the CPC to another, with a single instruction. This chip is used in building an easily reconfigurable, circuit switched connection network [Rana 88], for the interprocessor communication of the intermediate level processors of the Image Understanding Architecture (IUA) prototype [Weems 88].

1. INTRODUCTION

The interconnection structure between the processors of a parallel architecture greatly influences the run time performance of many parallel algorithms on the target system [Levitan 87, Gannon 84, Lint 81]. Different algorithms exhibit different communication requirements in terms of their communication patterns. Parallel systems with a static interconnection topology have many limitations [Lee 88] when they are used to run algorithms with different communication patterns. A parallel system whose interconnection topology can be configured to match the communication characteristics of an algorithm, can often achieve a significant speed increase over a system with a static topology.

Machine vision is one of the most computationally intractable problems, and requires a very broad spectrum of techniques and algorithms from signal processing to knowledge-based symbolic computing. As a part of the research on machine vision at the University of Massachusetts [Hanson 87], our group is working on the development of a massively parallel architecture called the Image Understanding Architecture (IUA).

The organization of this paper is as follows. Section 2 presents a brief overview of the IUA. In section 3 we briefly present the interconnection network for the intermediate level processors of the IUA. The PARCOS chip is the main building block of this network. The architecture of the PARCOS chip is described in section 4. The connection pattern setup and reswitching scheme for the PARCOS chip, and how it effects the overall organization of the whole network are also discussed. The paper ends with a discussion of future work and conclusions.

2. THE IMAGE UNDERSTANDING ARCHITECTURE

This section has a brief overview of the Image Understanding Architecture (IUA). A detailed discussion of the IUA can be found in [Weems 88].

The Image Understanding Architecture is a massively parallel, multi-level system for supporting real-time image understanding applications and research in knowledge-based computer vision. The IUA integrates parallel processors operating simultaneously at three levels of computational granularity in a tightly-coupled architecture. Each level of the

IUA is a parallel processor that is distinctly different from the other two levels, in order to best meet the processing needs at each of the corresponding levels of abstraction in the image interpretation process. Communication between levels takes place via parallel data and control paths. The processing elements within each level can also communicate with each other in parallel, via a different mechanism at each level that is designed to meet the specific communication needs of each level of abstraction. A block diagram of the IUA system is shown in figure 1.

The low-level, called the Content Addressable Array Parallel Processor (CAAPP), is a 512×512 array of bit-serial processors designed to operate on arrays of pixels and to construct intermediate-level tokens from events in an image. At the intermediate level, a collection of 4096 16-bit processors, called the Intermediate Communications and Associative Processor (ICAP), is used for retrieving, comparing, and matching tokens, computing geometric relationships between tokens, and constructing new tokens that describe more abstract entities. At the high level, called the Symbolic Processing Array (SPA), a set of 64 processors capable of executing LISP programs supports computation involving inference, hypothesis generation and verification, analysis of uncertainty, model-based processing, and indirect control of processing at the lower levels. The processors at the CAAPP and the ICAP levels are controlled by a dedicated Array Control Unit (ACU) that takes its directions from the SPA level. Data is moved between adjacent processing levels via dual-ported shared memory layers.

A proof-of-concept prototype of 1/64th of the IUA is currently being constructed by the University of Massachusetts and Hughes Research Laboratories.

3. THE INTERMEDIATE LEVEL

The ICAP level of the IUA is built out of fast digital signal processing (DSP) chips (Texas Instruments TMS320C25). Each DSP chip has one serial input port and one serial output port, each of which is capable of a 5M-bit/sec data rate. These serial ports provide the basis for interprocessor communication within the ICAP and as such they form the set

of data sources and sinks that are linked by the network described here.

3.1 INTERCONNECTION NETWORK

The ICAP connection network is used to set up a connection pattern between the N output ports of the processors and the N input ports of these same processors. The connection network can be programmed on-line, to make a direct link from the output port of any processor to the input port of one or more processors. The PARCOS chip is capable of broadcasting, allowing the connection network to realize any of the possible N^N mappings of its input ports onto its output ports. All of the processors can send and receive data on their links at the same time. These links can be changed by the ACU at any time.

The 64-input, 64-output connection network for the IUA prototype uses 2 stages of 32×32 PARCOS chips. The PARCOS chips are connected to make a 64×64 crossbar switch with broadcast capability as shown in fig. 2. A detailed discussion of the network can be found in [Rana 88].

4. THE PARALLEL COMMUNICATIONS SWITCH

The PARCOS chip consists of a communication matrix of 32 bit serial inputs and 32 bit serial outputs, a control memory, a set of registers and associated read/write circuitry. The PARCOS chip organization is shown in figure 3. Multiple PARCOS chips can be used to build larger connection networks, such as the 64×64 network in the IUA prototype.

The communication matrix of PARCOS consists of 32 tree-structured multiplexers, each of which is a 1 of 32 multiplexer. All 32 input lines are connected in parallel to each of the 32 multiplexers. With this architecture, multiple outputs can be connected to the same input, providing *broadcast* mode capability. Figure 4 illustrates one multiplexer tree. It will be noted that there are two multiplexer trees, one made out of n-channel transistors and the other made out of p-channel transistors, with their outputs connected together. By properly sizing the two types of transistors, we have achieved near equal delays for both low-to-high and high-to-low transitions at the output. For any multiplexer, path selection

at any level of the tree is done with a single bit of a control word. Thus, 5 control bits are required to select 1 of 32 inputs for each multiplexer, or $32 \times 5 = 160$ bits for configuring the entire communication matrix.

The PARCOS control memory consists of 32 control words, where each control word contains the 32 bytes of 5 bits required for one configuration. The on-chip control memory is therefore constructed so that PARCOS can hold up to 32 of the most frequently used connection patterns for larger networks built out of this chip. The control memory is called the Connection Pattern Cache (CPC), because it is analogous to storing the most frequently used pages in a memory system cache.

The connectivity information for the communication matrix is stored serially into the control words. To write connectivity information in a control word of the CPC, first a row number is set in the Row Select Register (RSR). RSR is mapped into the chip's memory space, allowing the address bus in PARCOS to select the register, and the binary value on the data lines determines the row number. Next, 32 5-bit bytes are written into addresses 0 - 31. The memory location's address is the output port number and its contents determine which input port it is connected to. If only a subset of links need to be modified, this can be done by selectively writing only into locations corresponding to those links.

Reswitching the configuration of the communication matrix from one stored connection pattern in a control word to another requires a single write instruction, where the address of a new control word is placed in the RSR, and the control word's contents are loaded into the Control Pattern Register (CPR), activating a new connection pattern. Notice that the CPR allows a control word to be modified in the CPC without disturbing an existing configuration in the communication matrix. In many cases this feature allows the time to write a new connection pattern from the ACU into the CPC to be hidden while the processors are working on an algorithm.

PARCOS is implemented on a single 84 pin, 50,000 device, VLSI chip. It is a full custom design, built out of a 2 micron, P-Well, double metal, scalable CMOS technology available through MOSIS. Each CPC memory bit is a 6 transistor static RAM cell. The worst case delay in broadcast mode from one input to 32 outputs is less than 50nS. A

microphotograph of the chip is shown in figure 5.

5. FUTURE MODIFICATIONS:

The design of the PARCOS chip was limited by the number of pins and not by the silicon area. A study for redesigning the PARCOS chip is underway with the goal of providing a 64×64 communication matrix with more than one hundred control words. Also, mechanisms will be provided for self routing in these chips in $\Theta(1)$ time. A 64×64 single chip implementation will allow us to build a 4096-input, 4096-output network by connecting 512 of these chips in a modified 4-stage, strictly non-blocking Clos' [Clos 53] topology or by connecting only 192 of these chips in a 3-stage, rearrangeably non-blocking Benes' [Benes 62] topology.

Currently it is not possible to directly copy or combine connection patterns in the CPC. Adding the ability to copy one CPC connection pattern into another under the control of a mask register will allow us to build new connection patterns from old ones, which will reduce the time required to create a new pattern in certain cases.

6. CONCLUSIONS

The architecture of a custom VLSI parallel communications switch was described. One scheme for constructing a 64×64 network, built with these chips, was described and some options for building much larger networks were mentioned.

REFERENCES:

[Benes 62] Benes, V.E., "On Rearrangeable Three-Stage Connecting Networks", *Bell Systems Tech. Journal*, vol. XLI, no. 5, Sept. 1962, pp 1481-1492.

[Clos 53] Clos, C., "A Study of Non-Blocking Switching Networks ", *Bell Systems Tech. Journal*, vol. 32, no. 2, March 1953,pp 406-424.

[Gannon 84] Gannon, D.G, and Rosendale, J.V., "On the impact of communication complexity on the design of parallel numeric algorithms", *IEEE Trans on computers*, Dec 1984, pp 1180 - 1194.

- [Hanson 87] Hanson, A.R., and Riseman, E.M., " The VISIONS Image Understanding System.", *COINS Technical Report, University of Massachusetts at Amherst*. 1987.
- [Lee 88] Lee, I., and Smitley, D., "A synthesis algorithm for reconfigurable interconnection networks", *IEEE Transactions on computers*, June 1988, pp 691 - 699.
- [Levitan 87] Levitan, S.P., "Measuring communication structures in parallel architectures and algorithms", in: *The characteristics of parallel algorithms, L. Jamieson et al (Eds), MIT press, Cambridge, 1987*. Also see "Parallel algorithms and architectures: A programmer's perspective", Ph.D. dissertation, *COINS tech rep 84-11, University of Massachusetts at Amherst*, May 1984.
- [Lint 81] Lint, B. and Agerwala, T., "Communication issues in the design and analysis of parallel algorithms", *IEEE Transactions on computers*, Mar 1981, pp 174 - 188.
- [Rana 88] Rana, D., Weems, C.C., and Levitan, S.P., "An easily reconfigurable circuit switched connection network", *Proc 1988 IEEE Int Symp on Circuits and Syst*, June 1988, pp 247 - 250.
- [Weems 88] Weems, C.C., Levitan, S.P., Hanson, A.R., Riseman, E.M., Shu, D.B., and Nash, J.G., "The Image Understanding Architecture", *International Journal on Computer Vision, Vol 2, Number 3*, December, 1988.

The Image Understanding Architecture

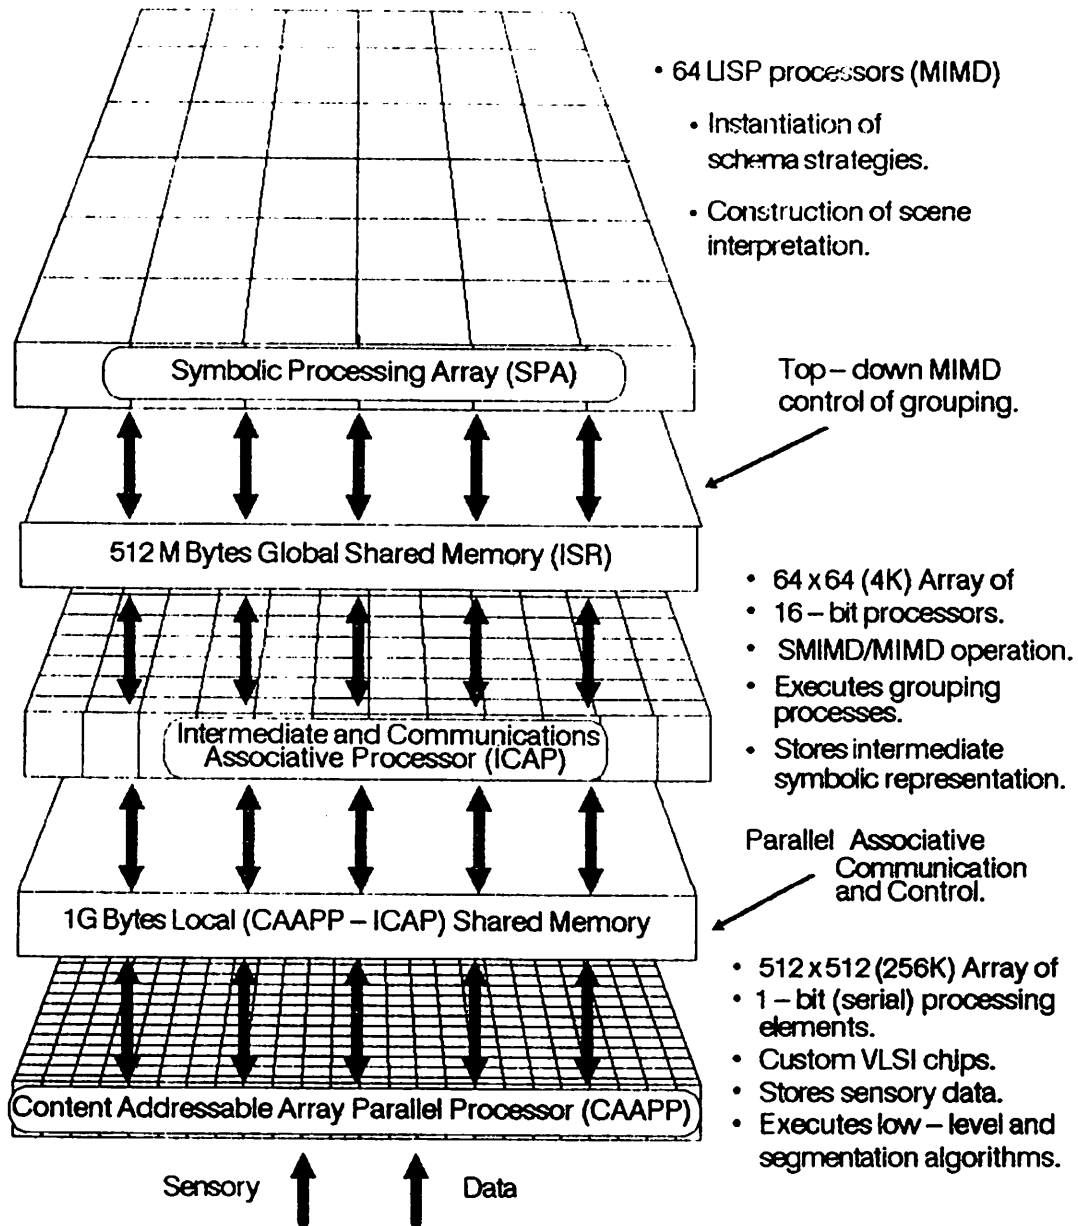



Fig. 1: UMass IUA Overview

Fig. 2: A 64-Input, 64-Output connection network



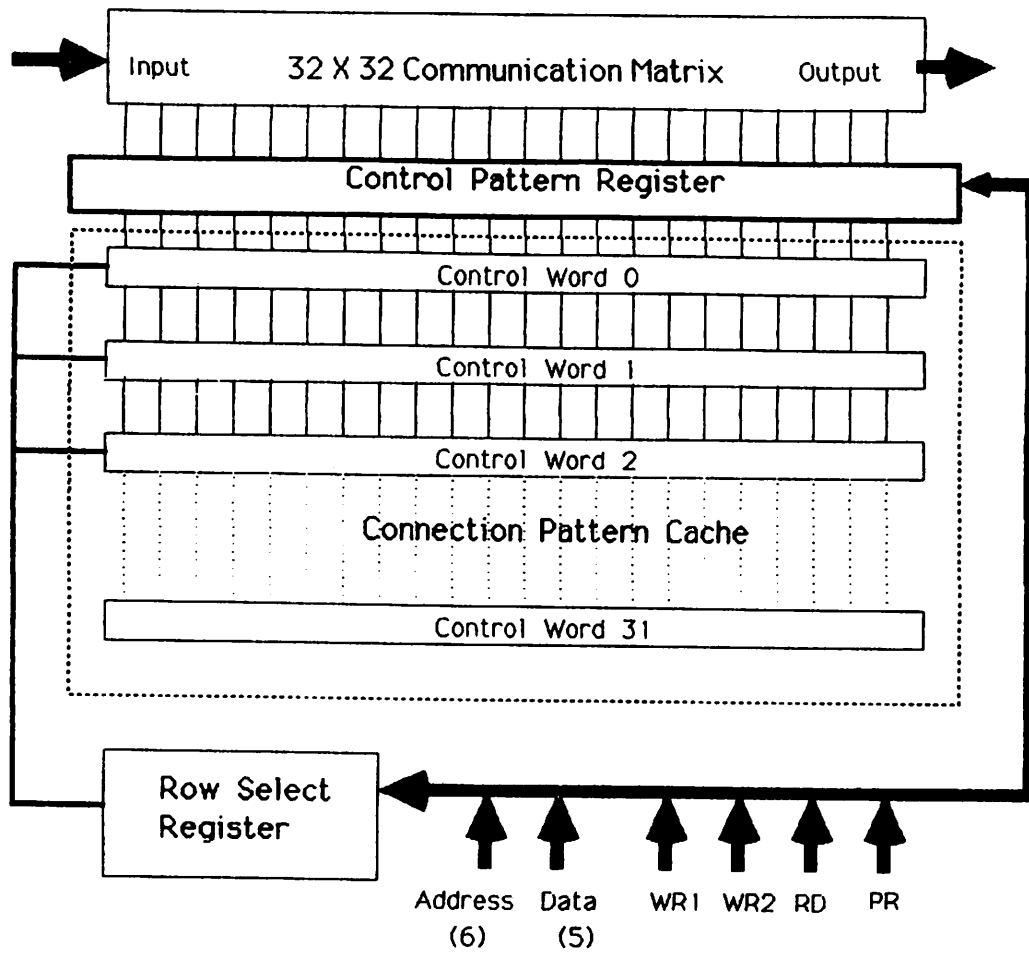


Fig. 3: Parallel Communications Switch Organization




Fig. 4: Multiplexer Tree




Fig. 5: Microphotograph of the PARCOS chip