ATCURE: A HETEROGENEOUS PROCESSOR FOR IMAGE RECOGNITION *

J. G. Ackenhusen, Q. A. Holmes, P. A. Kortesoja, D. L. McCubbrey, P. L. Mohan,
J. A. Salinger, T. N. Wessling, L. J. Witter

Environmental Research Institute of Michigan
P.O. Box 134001
Ann Arbor, Michigan 48113-4001

H. Stopper
Pico Systems
2901 Hubbard Drive
Ann Arbor, Michigan 48105

ABSTRACT

A processor that applies a heterogeneous architecture to the problem of real-time image recognition has been developed. Several unique features distinguish this work from other work in this field [1-4] and are the subject of this paper: 1) use of a *complete set* of documented applications of automatic target recognition to derive and validate processor requirements; 2) choice of a *heterogeneous* architecture that integrates several types of processors; 3) development of *image-processing-domain custom integrated* circuits; 4) application of *wafer-scale multichip module miniaturization* to the image processing pipeline; and 5) use of a *piecewise-connected hierarchy* of simulation tools, providing for connectivity of simulation both vertically (i.e., from chip through boards to subsystem) and horizontally (i.e., board vs. multichip module domain). This processor has been programmed with several recognition algorithms and delivered in a baseline 20-stage-pipeline-configuration, where it achieves over 20 billion Reduced Instruction Set (RISC)-equivalent operations/sec upon 16-bit pixels.

1. INTRODUCTION

Advanced image recognition requires multiple forms of computing, beginning with two-dimensional nearest-neighbor fixed point operations on pixels, progressing to the traditional multiply-accumulate floating point operations typical of time-domain signal processing, and concluding with the data-dependent symbolic-level processing required by evidential reasoning. Each domain is distinguished by different patterns of memory and operand usage as well as by differences in computational regularity, which means that each domain is best served by a different computational architecture (Fig. 1).

Past solutions to image recognition computing have often used a single architecture suited to one domain and have inefficiently extended its computation into the remaining domains. As an example, a digital signal processor has often been used in image recognition processing, where it has been forced to take on the necessary two-dimensional image operations (upstream) and the data-dependent decision operations (downstream).

The heterogeneous computer architecture, which combines a team of specialist processors and focuses their actions upon a single problem, has received recent attention as a solution to
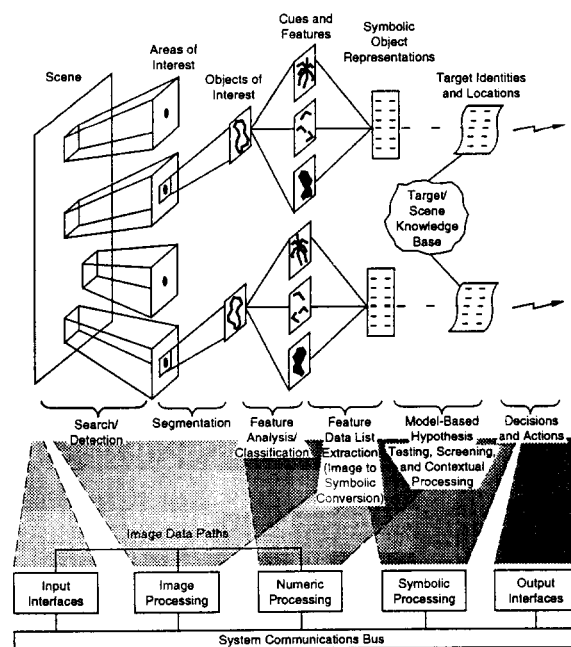

Fig. 1. Stages of generic image recognition flow diagram (top) map directly to components of heterogeneous architecture (bottom).

achieving the efficiencies of special-purpose architectures in the face of the breadth of real-world problems [5].

2. ARCHITECTURE DEFINITION

A processor that applies a heterogeneous architecture to the problem of real-time image recognition has been developed. The processor serves the domain of generic image interpretation applications, originating from the needs of automatic target cueing and recognition. This processor, known as ATCURE (for Advanced Target CUeing and Recognition Engine), was defined by first analyzing the complete set of future systems documented in the *U.S. Army Next Generation and Future Systems Sourcebook* [6] and selecting the 39 system areas that required capabilities in automatic target cueing or recognition (Fig. 2).

The concept of a *mission snapshot* was devised to capture the time-critical parameters that pace the real-time execution of image recognition algorithms. Such parameters include frame

size (length and width in pixels), frame rate (in frames/sec), pixel depth (in bits per pixel), expected number of targets in a scene, rate of platform movement, and response time. Representative algorithms from each system were thoroughly analyzed. The computational burden of each snapshot from each system was assessed for various computational architectures, and an architectural simulator was used to refine estimates of bus timing. Iterations were made between architectural alternatives and performance prediction, resulting in the choice of a heterogeneous architecture.

of 50 to 70 percent are typical in image processing. A baseline subsystem configuration uses a pipeline of 20 PPEs, and software provides a virtual pipeline that allows its effective length to be extended indefinitely through recirculation.

The IPS image processing throughput is maximized by an intelligent image memory/region controller custom integrated circuit, which provides hardware acceleration of the complicated two-dimensional addressing that is typical of image operations (e.g. area-of-interest segmentation, up/down-sampling, corner turning, table lookup). A scaleable high-speed

Fig. 2. The requirements of ATCURE were developed from the future Army systems that will use automatic target recognition.

3. IMAGE PROCESSING SUBSYSTEM

The heterogeneous architecture maps directly to the generic flow of image recognition algorithms (Fig. 1). For the front end computations, image processing computations are executed by the Image Processing Subsystem (IPS). Central to the IPS design is the principle that image processing operations can be decomposed into sequences of nearest-neighbor pixel operations, scanned across the image. A raster pipeline subarray processor known as the Pipeline Processing Element (PPE) uses an application-specific integrated circuit to capture the computational kernel that simultaneously presents a pixel's set of nearest neighbors to an arithmetic/logical neighborhood processor. The IPS consists of a pipeline that combines multiple PPEs with 16 global image buses and several additional local (direct PPE-to-PPE) buses, each of which transfers 16-bit pixels at a 20 Mpixel/sec rate. Each PPE performs 20 million neighborhood operations/sec, where a neighborhood operation combines a pixel's nine nearest neighbors via simultaneous use of 10 multipliers, 29 adders, 2 ALUs, 2 counters, and an accumulator, and then outputs a pixel that is a linear or non-linear combination of its neighbors. A peak rate of 44 operations/neighborhood x 20 MHz = 880 million operations/sec per PPE may be achieved -- efficiencies

crossbar switch, implemented with a custom integrated circuit building block, supports the image buses. The IPS is programmed in IPSTran, a machine-specific language capable of producing C-callable modules. Throughput rates at 20 Mpixel/sec processing for several operations are shown for both a single PPE and a 20 PPE pipeline in Table 1.

Table 1. Typical Execution Times on 1000x1000 Pixel Image

Operation	Single PPE	20-PPE IPS	Reference Sun 4/330
3x3 convolution	0.05 sec	0.05 sec	6.1 sec
9x9 convolution	0.15 sec	0.05 sec	8.6 sec
Sobel edge detector	0.45 sec	0.05 sec	53 sec
16 stage morphological filter	0.80 sec	0.05 sec	496 sec

The Numeric Processing Subsystem follows the IPS and uses an array of TMS320C40 standard signal processors to perform the multiply-add-intensive floating point signal processing operations associated with geometric correction, scaling, and feature analysis. A standard microcomputer based on the Motorola 68040 microprocessor serves as the Symbolic Processing Subsystem to execute the data-dependent processing

2680

associated with recognition decision processing. Both the numeric and symbolic processor are programmed in C.
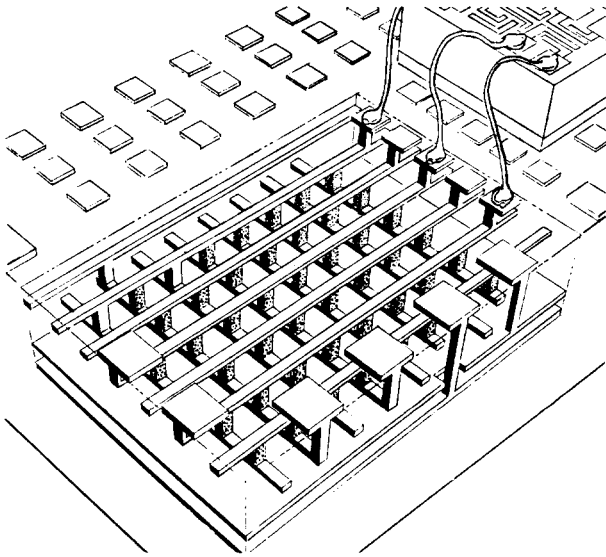

Fig. 3. Programmable silicon circuit board with antifuses.

4. MINIATURIZATION

The feasibility of miniaturization of the processor was demonstrated by miniaturizing the image processing pipeline. A novel approach to wafer-scale integration was used that employed customized programmable silicon circuit boards (PSCB) as an interconnecting substrate to wire-bonded bare silicon die [7]. As shown in Fig. 3, a PSCB consists of a layered silicon substrate that provides a level of parallel conductors running orthogonal to a lower level of parallel conductors. At each intersection is an electrically-programmable "antifuse" that can be permanently transitioned from an insulating state to a conductive state. Because design-specific interconnection information is electrically injected into the PSCBs, this technology supports the rapid-prototyping, small-run approach more easily than those multichip module technologies that require the design-specific interconnect to be imposed as mask layers during substrate fabrication (the analogy is similar to electrically-programmed vs. mask-programmed read-only memories).

As a result of this approach, it was possible to achieve a 4-inch-diameter multichip module, allowing an entire 9U VME board of conventional circuitry to be collapsed to one multichip module (Fig. 4). This preserved the circuit partitioning between the standard and the miniaturized version of the circuits. Using multichip module technology introduced design considerations regarding circuit speed and routing early in the design process [8] and resulted in the development of the concept of *incremental functional test*. Because the wafer-scale miniaturization technology allowed the construction of a high-value assembly containing several dozen integrated circuits, the traditional "build, then test" approach was replaced by the incremental "build a little, test a little" approach to accelerate the detection of defects and eliminate the chance of having to scrap an entire finished module.

5. SIMULATION AND TEST

The diversity of application requirements, the circuit parameters of multichip modules, and the integration demands of a heterogeneous configuration of diverse processors presented challenges to simulation of the processor throughout its design. Several factors contributed to rapid system integration and correct-first-time operation:
- modeling of the analog circuit parameters of the programmable silicon circuit board and using these in circuit design;
- a connected hierarchy of simulation provided by the software tools of the commercial chip vendor that allowed gate level simulation of the pipeline processor integrated
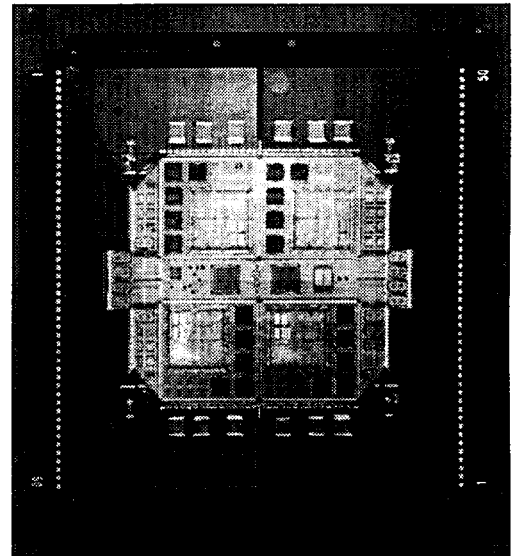

Fig. 4. Wafer-scale multichip module contains four PPEs.

circuit to be combined with simulation of the image memory custom integrated circuit and the other circuits on a board and the cascading of several boards into a subsystem in the simulation domain;
- use of "virtual prototyping" techniques in which actual image processing component algorithms were run on the above-mentioned hierarchical simulation model well in advance of actual hardware;
- use of identical circuit models and test fixtures on the non-miniaturized and miniaturized versions of the circuit, accelerating the introduction of miniaturized modules into the system.

As a result, the four custom printed circuit boards, three custom integrated circuits, 100,000-line system software, and miniaturized module were integrated within one month and were running applications within one more month.

6. CONCLUSIONS

The ATCURE processor has been programmed with three diverse end-to-end algorithms for image recognition, each using a different input sensor data type and seeking different objects, with one fusing input from multiple sensors (Table 2).

Table 2: Diverse Automatic Target Cueing and Recognition
Algorithms Implemented on ATCURE

Algorithm	Sensors	Comment
Critical Mobile Target	synthetic aperture radar (SAR)	distinguishes certain vehicles from clutter in fine-resolution SAR imagery
Background Adaptation Convexity Operator Region Extraction (BACORE)	forward-looking infrared (FLIR)	recognized as of complexity representative of emerging algorithms
Remote Minefield Detection System (REMIDS)	three sensor types (polarization, reflectance, thermal)	sensor fusion combines separate views of same region

Figure 5 displays the ATCURE hardware with a simulated synthetic aperture radar image as input and the resulting detection of vehicles of a specific type as output (in boxes) and their discrimination against other objects (not boxed). ATCURE has been delivered to the U.S. Army Night Vision and Electronic Sensors Directorate, where it is being used in various demonstrations.

Efforts to extend the ATCURE technology include:

- introduction of the IPS into existing open architecture computers in need of image processing functionality;
- development of workstation-embedded image processing accelerators using the components of the IPS;
- application of the programmable silicon circuit board rapid prototyping technology to other circuits;
- introduction of image recognition algorithms for character recognition, sighted automation, and image information extraction/fusion onto the ATCURE processing elements.

REFERENCES

[1] McCloud, Eugene L., "Geometric Parallel Processor: Architecture and Implementation," *Parallel Architectures and Algorithms for Image Understanding*, Academic Press, pp. 279-305, 1991.

[2] Branstetter, R., Roarke, C., and Ruszczyk, W. "ALADDIN Processor and Software Support." In *Digest of Proceedings Government Microcircuit Applications Conference*, Las Vegas, NV, pp. 493-496, 6-9 November 1990.

[3] Belt, R., *et al.*, "The ALADDIN Processor: A Miniaturized Target Recognition Processor With Multi-GFLOP Throughput." In *Proceedings of the SPIE Conference on Architecture, Hardware, Forward-Looking Infrared Issues in Automatic Target Recognition* (SPIE-1957), Orlando, FL, pp. 264-275, 12-13 April.

[4] Adams, C., "Chasing the Elusive Gigaflop in a Soup Can," *Military & Aerospace Electronics*, pp. 26-29, Feb. 1991.

[5] *Computer,* Special Issue on Heterogeneous Processing, June 1993.

[6] *Next Generation and Future Systems Notebook*, U.S. Army Materiel Command, 28 November 1989.

[7] Stopper, H., "An Advanced Version of the Electrically-Programmable Hybrid-WSI Substrate." In *Proceedings of the International Conference on Wafer Scale Integration*, pp. 289-298, Jan. 1993.

[8] Banker, Jeffrey and J. Muczynski, "Designing with Programmable Multichip Modules," *Proceedings of Electronic Packaging Conference of Packaging Exhibition,* San Diego, CA Sept. 1993.

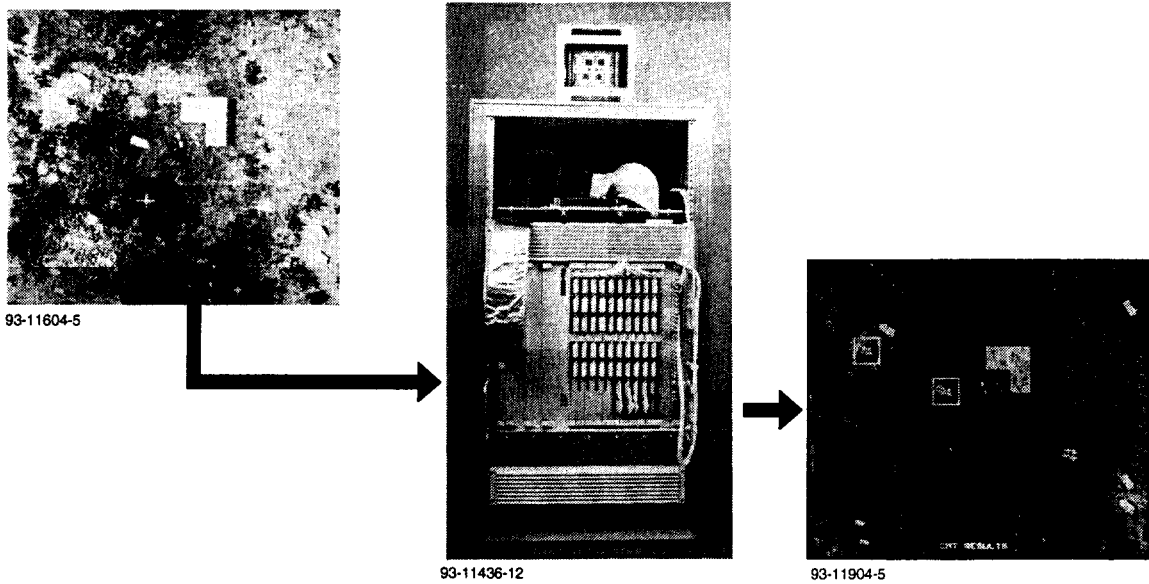93-11604-5

93-11436-12

93-11904-5

Fig. 5: ATCURE hardware with miniaturized pipeline (box on top) performs recognition on simulated SAR imagery.