CONCURRENT PROCESSOR ASIC FOR HIGH SPEED PATH PLANNING

S.E. Kemeny, T.J. Shaw, R.H. Nixon, T. Daud, and E.R. Fossum
Jet Propulsion Laboratory/California Institute of Technology
Pasadena, California 91109

Introduction

For a given terrain to be traversed, it is computationally intensive to determine the fastest route between two points, and for defense or civilian emergency dispatching applications, computation time is critical. This paper reports the integration of a 24 x 25 random access array of digital processors which are programmed to model a given terrain and determine the fastest (lowest cost) path from an origination point (or a multiple of points) to all the other points on the terrain at very high speed (milliseconds for arrays up to 512 x 512). The primary purpose of this research chip is to demonstrate high speed path planning capability for tactical mobility analysis in battlefield scenarios. However such a high-speed, automated path-planning processor will find utility in a variety of settings such as autonomous vehicle navigation, intelligent vehicle highway systems, evacuation and rescue planning, and police and transportation dispatching.

Currently, the only tools available to assist in path planning are implemented in software. These approaches are slow, with best-path determination typically requiring seconds to minutes for terrain sizes varying from 64 x 64 to 512 x 512 pixels.[1] Through the VLSI implementation of a fine grain parallel architecture, in which every terrain pixel is represented by a corresponding processor, the inherent parallelism of the problem can be exploited and extremely fast path determination can be realized. In such an architecture, the only processor communication required is between nearest neighbors so that processor communication overhead is virtually eliminated. This is in contrast to conventional parallel computers, where even with proper parallel decomposition of the problem, processor communication overhead is often a severe speed bottleneck. In this paper, the first parallel processor IC for route-planning over complex terrain is reported.

Array Architecture and Operation

The path planner architecture, shown schematically in Fig. 1, consists of a 24 x 25 array of unit cells (processors) which communicate with their nearest neighbors and are randomly accessed by 5-bit row and column decoders located adjacent to the array. The IC is implemented in a single-poly, double-metal 2 µm CMOS n-well process, utilizing a full custom layout. The unit cell occupies 296 µm x 330 µm and the overall chip area is 9.2 mm x 7.9 mm.

In order to determine the fastest routes from a selected starting point(s) to all other points on a given terrain, each unit cell corresponds to a terrain pixel which has been preprogrammed with the cost (i.e. delay) of traversing that pixel. Operation begins with the selection of a path origination pixel(s) which sends out a signal to its north, south, east, and west neighbors. Each neighbor delays the signal by a preset time (programmable cost), after which it broadcasts a signal to each of its four neighbors. When a signal is received, the incoming signal direction is stored and further inputs to the cell are disabled. This results in a signal wavefront propagating radially outward from the originating pixel that is then distorted by the varying delays encountered in the array. The variable unit cell delays are implemented by

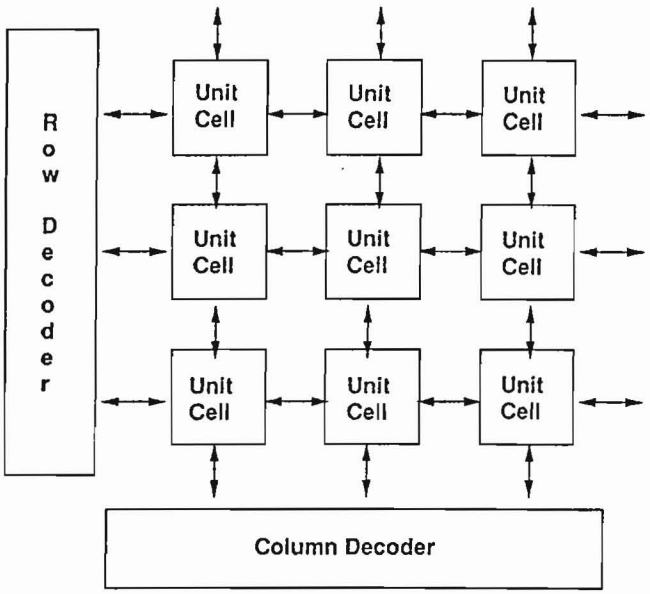

Figure 1. Block diagram of the IC architecture.

presetting an 8-bit ripple down-counter in each cell to one of its 256 possible values (costs). When triggered by an incoming signal, the counter decrements down to zero, which in turn triggers the broadcast of an outgoing signal to each of its nearest neighbors. When signal propagation through the entire array is complete, any destination node may be queried and the minimum path between it and the origination node is found by retracing the direction stored in each unit cell. Thus, determination of the fastest paths through a complex terrain (modeled by 256 cost levels) is realized. This is in contrast to the simpler task of maze solving or wire routing, in which the processors would be programmed with binary costs, i.e. the pixel is either blocked or open.[2]

In addition to finding the fastest paths from one origination pixel to all possible destinations, multiple starting pixels can be selected, with signal propagation emanating from each source and stopping at the boundary between signal wavefronts. This feature is useful in battlefield scenarios where an analyst can model the progress of different forces across the terrain. In addition, when any destination node is queried, the minimum path between it and the nearest source pixel is displayed, which provides valuable information for rescue operations.

Another unit cell function is the conditional blocking of signal propagation in any direction to model impassable terrain such as rivers and canyons. This is important because the current (and next generation) resolution of digitized map data results in single pixels which contain both rivers and other features such as roads. Therefore, it is important to be able to assign the terrain pixel the cost of a road, while blocking the signal from crossing the river. In this way signal propagation can occur along a road adjacent to a river. Such

conditional blocking is accomplished with a set of four static latches which are preprogrammed to either block or transmit the signal emanating from the counter.

Experimental Results

The path planner chip was interfaced to a laboratory PC computer through a wirewrap board and plug-in digital interface card. The chip (address memory, counter, river blocking, control logic and I/O) is completely functional. A summary of the chip characteristics is given in Table 1. It was found that the latches require 160 ns to settle, implying a terrain programming time for the 600 pixel array of less than 300 μs. A separate counter test circuit was successfully clocked at 8.33 MHz, limited by the test station. The array can be operated in two modes: single step and continuous. In the former, the chip is clocked via the PC, and the actual signal propagation on the chip can be monitored on the PC screen. In the continuous mode, a function generator supplies a square wave input to a non overlapping clock generator located on the wire-wrap board which in turn clocks the counter. The chip was tested at frequencies up to 7 MHz in this mode, resulting in typical path determination times of under 250 μs. For a typical terrain cost map, signal propagation through the array required 2550 clock cycles, so that the entire signal propagation phase required only 360 μs at a 7 MHz clock rate. Fig. 2 displays the original map with a typical lowest cost path shown in white.

Table 1. IC Characteristics

Chip Architecture	24x25 digital processor array
Maximum Clock Frequency	7 MHz
Programming Speed	2 Megapixels/second
Equivalent Operations	6 billion/second
Origination Nodes	One or Multiple
Dynamic Range of Cost	256:1
Chip Fabrication Process	2 μm CMOS
Unit Cell (Processor) Size	296 μm x 330 μm
IC Size	7.9 mm x 9.2 mm

Conclusion

In summary, the first single-chip fine grain parallel processor array to perform path planning over complex terrain has been demonstrated. The 24 x 25 array of digital processors has been operated at frequencies up to 7 MHz, providing best (fastest) route determination in under a millisecond. This corresponds to a four orders of magnitude speed-up over current software approaches. Full functionality of this first generation research chip paves the way for the implementation of large arrays (e.g. 1024 x 1024) and chips with increased functionality. Both these avenues are currently being pursued.

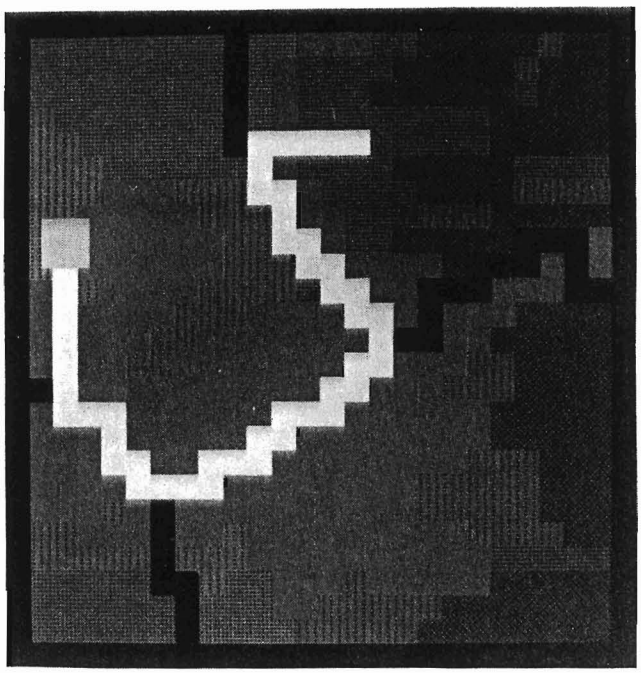

Figure 2. A typical lowest cost path found by chip.

Acknowledgments

The authors gratefully acknowledge useful technical discussions with H. Langenbacher, T. Brown, S. Eberhardt, and A. Thakoor during the course of this work. Special thanks go to D. Kerns for his effort on a previous analog path planner design and to B. Minch who wrote the original software interface for operation of the path planner chip. The research described in this paper was performed at the Center for Space Microelectronics Technology, Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration and was jointly sponsored by the All Source Analysis System Program Office, the Defense Advanced Research Projects Agency, and the Naval Surface Warfare Center.

References

1. T. Kreitzberg, T. Barragy, and N. Bryant, "Tactical Movement Analyzer: A Battlefield Mobility Tool," Proc. Joint Service Data Fusion Symposium, Laurel, MD, 1990.

2. C.R. Carroll, "A Neural Processor for Maze Solving," in C. Mead and M. Ismail, Eds., Analog VLSI Implementation of Neural Systems, Kluwer Academic Publishers, Boston (1989).