Real-time processor for staring receivers

Brian Hanzal, Andrzej Peczalski, James Schwanebeck,
Richard Sanderson*, Eric Fossum**

Honeywell Systems & Research Center, Minneapolis, MN
* Wright Laboratory, WPAFB, Dayton, OH
** Jet Propulsion Laboratory, Pasadena, CA

ABSTRACT

The design, fabrication, and testing of a state-of-the-art, high-throughput on-focal plane infrared-image signal processor is described. The processing functions performed are frame differencing and thresholding. The final focal plane array will consist of a 128x128-pixel platinum-silicide detector bump-mounted to an on-chip CCD multiplexer. The processor is in a 128-channel parallel-pipeline format. Each channel consists of a pixel regenerator (charge differencer), 128-pixel frame store CCD memory, pixel differencer, second pixel regenerator, thresholder (analog comparator), and digital latch. Four parallel analog outputs and four parallel digital outputs are included. The digital outputs provide a bit map of the image. All analog clock signals (128 KHz, 256 KHz, and 5 MHz) are generated by on-chip TTL-input clock drivers. TTL clock driver inputs are generated off-chip. The technology is low-temperature surface and buried channel CCD/CMOS/indium bump. It is believed that this is the first time that all of these technologies are integrated into one device. The design goal was 8-bit resolution at 77K and 1000 frames per second. The chip is 890x900 mils^2. Applications include point- or extended-target motion detection with thresholding. Design trade-offs and enhancements (such as on-chip detector gain compensation and a simple window processor) are discussed.

1. INTRODUCTION

The signal processing required for imaging systems (pixel compensation, signal-to-noise enhancements) demands very high throughput on the order of billions of operations per second. Thus specialized, high throughput integrated circuits are desired. The CCD-based processor offers high parallelism (in our case 128 processing channels perform all of the algorithms in parallel). The throughput-of-power figure of merit of such a processor is several times better than any digital preprocessor. The second advantage of the preprocessor is that all of the processing is performed before analog-to-digital conversion. Thus the data rate can be reduced before it arrives at the analog-to-digital converter, which reduces the system power and size. The third advantage of the CCD processor is that the processing occurs at the focal plane, which is either integrated with or bumped to the processor and therefore simplifies the interconnection between the sensor and the rest of the system.

The remainder of the paper is divided into nine sections. They will present, in order, a brief comparison of analog and digital approaches, a chip overview, a detailed description of the sections of the processor, the test circuits, the technology, experimental results, and proposed enhancements. The paper ends with acknowledgements and references.

2. SIZE, POWER TRADEOFF BETWEEN RTP AND DIGITAL PROCESSOR

Since such systems are also being implemented in digital technology, tradeoffs with the Real-Time Processor (RTP) performance are required. During the first part of the program we performed a quick and simple tradeoff. We used as an example of a digital image signal processor a 16-bit parallel-recirculating-pipeline (PREP) processor also designed at Honeywell.

The data flow and functions of a proposed digital implementation of the RTP is shown in Figure 1. The 128 X 128 pixel sensor array running at 1000 frames per second is assumed. This results in a 16.4 MHz pixel rate. The first PREP module performs gain compensation using the gain coefficients stored in a 16K x 8 bit SRAM. This data is then stored in a frame buffer memory which is another 16K x 8 SRAM. The frame differencing is performed in another PREP chip. The most difficult operation is the Laplacian filter in the 3 x 3 window, since the PREP chip operates only on 2 x 2 windows. Because of this, 11 pipelined PREPs are needed for the Laplacian window operator. The last of the 11 PREPs can also perform the thresholding operation. The output of the threshold operation is the bit map. A total of 17 chips (one circuit board) and 26.5 watts are necessary for the digital implementation of an RTP equivalent. In contrast, the RTP device is one integrated circuit dissipating approximately 2 watts. The power comparison advantage of RTP becomes even more evident when configuring a similar system for a 256 x 256 detector array. In order to match the enhanced performance of the RTP, it is estimated that 60 integrated circuits and 118 watts of power are required.

It was also found that above this number of pixels, and for more complex algorithms, the digital approach appears to be more advantageous. The advantage of the digital implementation is the flexibility of the algorithms. The precision of the digital system is limited by the precision of the analog-to-digital converters.

This tradeoff compared a specific analog processor for a specific application with a specific digital processor which was not necessarily optimized for that application. Other scenarios and processors may give other results. Such exhaustive tradeoffs were not within the scope of this project.

3. CHIP OVERVIEW

Typical imaging system requirements indicate at least an eight bit precision processor. This chip was originally intended for point targets but in its present configuration it has capability for extended targets.

A block diagram of one of the 128 channels is shown in Figure 2. In its final configuration the device will consist of a 128x128 hybrid platinum silicide imager bump-interconnected to a much larger main chip containing a multiplexer and a preprocessor. However, it is presently being tested as a visible imager using on-chip photodiode structures. The multiplexer is a 128x128 matrix of CCD readout cells. The imager cells are 50 microns center-to-center. There are 128 pipelined preprocessors ('channels') arranged in one row underneath the multiplexer. The processor channels are 142 microns center-to-center. The multiplexer and channel processors are interconnected via a set of 128 fanout CCDs. The processors are connected to a readout structure via a set of 128 fanin CCDs. The chip has capability for both four analog and four digital (bit map) outputs. Each output serves 32 multiplexer/processor columns.

A plot of the first metal layer of the chip is shown in Figure 3. The width and height of the chip is 22600 microns (889.8 mils) and 22850 microns (899.6 mils). The charge handling capability is limited by the smallest gate area in the multiplexer. This gate area is 470.25 micron^2 giving a charge handling capacity of near 940,500 electrons.

The design goal was 1000 frames per second. This implies 128,000 rows per second (7.81 μsec). The overall data rate is then 128x128x1000 = 16.384 MHz. Since this is divided among four parallel outputs, the maximum data rate from one output is 4.1 MHz. A typical clock swing is from 10 V to 0 V.

Because of all of the CCDs used to transport charge in the chip, it will take a certain amount of time for data to appear at the output once it is clocked out of the multiplexer. From the time a pixel is clocked out of the bottom row of the imager multiplexer until the time it enters the readout multiplexer it will have been stored in approximately 352 CCD stages (not counting any storage due strictly to the processing operations and the frame store CCD itself). This results in 2.75 frames stored in the processor at one time and latency of 2.75 frame periods. We do not count the charges in the serpentine memory because once it is filled with the first frame, it does not contribute to a delay in the output.

4. DETAILED DESCRIPTION

This section presents a discussion of the different elements of the Real-Time Processor.

4.1. Detector chip

The detector chip to be used is a backside-illuminated, high fill factor 128x128 platinum silicide (Schottky-barrier) detector array. The detectors/bumps are on 50-micron centers.

4.2. Multiplexer

The multiplexer has three-phase vertical shift registers. The reason for this is that it is not possible with our 50 micron cell to get another clock line into the cell for the fourth phase. The multiplexer gate areas were made as large and equal as possible. The smallest of these is 470.25 micron^2 and will provide the upper limit for the processor's charge handling capability. The multiplexer has a three-phase horizontal readout shift register. This shift register also has the capability of passing charge packets vertically through it to the fanout CCD which transports the charge to the processor channel.

In order to implement the vertical pass-through feature and still have the horizontal multiplexer be capable of handling the full charge would have required a pass-through gate length of 52 microns. Since the vertical pass-through feature is the main mode of chip operation, we made this gate as short as possible at the expense of the area of the other gates in the horizontal shift register. The other gates in the horizontal CCD (used for testing only) are smaller than 470.25 micron^2 and will not handle the full charge.

Both the vertical and the horizontal shift registers have fat zero inputs for testing.

4.3. Output amplifiers

The output amplifiers consist of two source follower stages separated by a sample/hold MOSFET. It is preceeded by a floating diffusion on the output of the CCD.

4.4. Clock drivers

CCDs generally require clock voltage swings which are significantly larger than those provided by digital TTL or CMOS drivers. Diodes generally need voltage swings which are of different levels than those required for the CCD gates themselves. In addition, these waveforms must have rise and fall times slow enough to allow the smooth passage of charge. This in turn requires a discrete analog waveform generator for each unique clock in the system. By using on-chip clock drivers, we reduced the problem to that of supplying a digital TTL signal for each CCD clock. The non-digital voltage swing required by the CCD gates and diodes is provided by connecting each clock driver to a high and a low DC rail. In our case there are eight rails, corresponding to the high and low voltages required for diodes, first polysilicon gates, second polysilicon gates, and gates consisting of adjacent first and second polysilicon. These rails are implemented as second metal 'fences' around the main structure.

The clock driver design consists of a TLL-to-CMOS level-shifter front end and a CMOS-to-CCD voltage backend. The capacitive loading of each line was determined. The drive capability of the associated clock driver was then tailored to give the required rise and fall time.

4.5. Multiplexer-to-processor fanout fold-over CCD

The multiplexer cells are on 50-micron centers and the processor channels are on 142-micron centers. Thus, a straight CCD from a multiplexer column to its associated processor is not possible if we require that each charge sees the same number of approximately equal area gates. Since the CCDs must be of equal length and the center processors are closer to the multiplexer than the outer processors, a CCD with two-folds was used. One column of the fanout CCD is made up of a single CCD folded back on itself. As we move from the center out, the distance by which a CCD folds back on itself decreases to take up the added distance between a multiplexer column and its associated processor. It is called the fanout CCD since it fans-out from the small multiplexer to the processor. The number of these gates makes a contact to each one fairly risky. Some of the gates are too small to accommodate a contact. Therefore, the multiplexer-to-processor fanout CCD uses the 'sea-of-poly' concept. Many lines of polysilicon traverse the entire chip and are connected to clock lines at the chip edges only. The assumtion in this approach is that the CCD is clocked slow enough (≈128 KHz) that the clock driver can drive the higher capacitance and sheet resistance ('RC' constant) of the polysilicon so that the channels at the edge and the middle of the chip see the same gate voltage during the cycle.

4.6. Processor (from the first regenerator to the latch)

This section discusses the channel processor portion of the chip. It describes the operation of one of the 128 parallel, identical processor columns. Therefore, when we say that the processor does something to one charge packet, in practice one entire row of the 128x128 scene is being processed simultaneously by the whole chip.

The center-to-center spacing of the processor channels is 142 microns. Thus the width of the processing area is 128 x 142 = 18176 microns, plus a short distance in the middle because the left side is a mirror reflection of the right side.

4.6.1. First regenerator

The regenerator creates two copies of the original pixel value (which is lost in the regeneration process). One copy goes into the serpentine memory (to be used as the previous frame value) and the second copy is used immediately to perform a frame difference (using the corresponding pixel from the previous frame which has been stored in the serpentine memory). The transfer function of the differencer/regenerator depends on the parasitic capacitance of the principal gates. This in turn depends to a great extent on the overlap capacitance between the first polysilicon principal gate and the second polysilicon 'compensator' gate. The analysis of the regenerator follows that of Fossum[1] and Eid[2].

The variation in oxide capacitance (per unit area) due to variations in the oxide thickness both for polysilicon-to-channel and first polysilicon-to-second polysilicon overlap is just a few percent. The variation in total first polysilicon-to-second polysilicon overlap capacitance will involve both the variation in oxide thickness and variation in area due to misalignment and rotation. The misalignment variation could be up to ±0.5 micron. Thus, for any given overlap area, a misalignment or rotation will severely alter the overlap capacitance. However, since both of the first polysilicon principal gates overlap second polysilicon gates on both sides in the differencer, any misalignment or rotation should result in a zero change in total overlap capacitance.

4.6.2. Serpentine (frame store) memory
The serpentine memory is constructed of the 'sea-of-poly' concept discussed above. It stores one column of data from the multiplexer (128 values). It is a U-shaped stack which receives one of the outputs from the first regenerator and at the same time gives up the corresponding pixel from the previous frame to the frame differencer.

4.6.3. Differencer
The differencer performs a subtraction of analog charge packet values to create a frame difference. The frame difference will show a high output pixel only on the first frame in which the target enters the pixel's field of view. The regenerators are merely differencers in which one of the input charges is zero. The analysis of the differencer and the regenerator follows that of Fossum[1] and Eid[2].

4.6.4. Differencer-to-second-regenerator CCD
There is a CCD from the output of the frame differencer to a second regenerator. It runs parallel to the serpentine memory and is incorporated in the serpentine memory 'sea of poly' structure. The frame difference operation is carried out at the upper part of the channel. This is done so that the current pixel is available to the differencer;. This CCD is required in order to transport the pixel difference from the top of the channel down past the frame store memory to the second regenerator.

4.6.5. Second regenerator
This creates two copies of the frame difference. One copy goes to the comparator to be compared with a threshold charge packet and the other goes to the analog output multiplexer.

4.6.6. Comparator
The comparator is a charge comparison circuit. Its simulation equivalent circuit is shown in Figure 4. In the actual implementation, the current source inputs are replaced by CCDs and the C and E FETs are just overlapping gates with no intervening diffusion. Its operation is decribed by Eid[2] or Fossum[3]. The comparator works by charging and isolating two nodes with different quantities of charge (data and threshold). Each node has a gate which controls the discharge of charge on the other node. Thus the node with the largest amount of charge will allow the other node to discharge first, pulling one side high and forcing the other side low. In particular, the nodes must present equal capacitive loads to the two incoming charge packets. The layout must be performed not only so that the capacitance is equal but also so that it is equal in case the two polysilicon gate layers are misaligned to themselves or to field cut. Our design allows for a misalignment of up to 1.5 microns without a capacitance imbalance.

4.6.7. Latch
The latch is a modified version of a CMOS D-latch layout from Honeywell. A MOSFET pass gate has been inserted in the Q output path. This pass gate is clocked OFF when data is coming into the latch from the comparator and then clocked ON for serial latch readout. There are two versions of the latch, one with the normal internal enable signals (E, /E) and one with these signals interchanged to allow the use of only one external clock rather than two clocks 180 degrees out of phase. The 128 latches are grouped into four strings of 32 each. The last latch of each group goes to a CMOS output buffer. The digital output buffer design was suggested by Loral.

4.7. Processor-to-output fanin fold-over CCD
The processor-to-output-multiplexer fanout CCD is constructed of the 'sea-of-poly' concept discussed above. It is called a fanin CCD because it fans-in from the wide processor to the 50-micron stages of the readout multiplexers. It was originally intended that the analog output of the chip consist only of the magnitudes of the first sixty-four pixels which were above threshold. However, no acceptable implementation of such a feature was found. Using four fanin CCDs (a smaller, vertically-flipped version of the fanout CCD) and four output CCDs, we now read each pixel value and save only the sixty-four over threshold off-chip.

4.8. Output multiplexer CCD
The fanin CCD feeds into a horizontal four phase readout CCD. This horizontal shift register has a fat zero input for testing.

5. TEST STRUCTURES

Referring to Figure 3, there are 15 test sites within the main chip, placed on either side of the imaging array. Some of these contain more than one test structure. Other test sites are located outside the main chips to more efficiently utilize real estate. In all there are 24 test structures. They include both fundamental elements used in the present baseline processor and devices to be

assessed for use in an enhanced version. There are also several test structures representing various subsections of the main procesor and combinations of basic elements. There are six basic elements whose performance will indicate the performance of the overall chip. They are: a sample clock driver, a latch/string of latches, the differencer, the comparator, the output amplifier, and a four-stage, four-phase CCD. In particular, there are five differencers to test variations in the overlap capacitance, and thus test the transfer function.

6. TECHNOLOGY/PROCESSING CHALLENGES

The design uses a two polysilicon/two metal/surface and buried channel CCD/CMOS/indium bump process optimized for low temperature operation. This required sixteen mask layers and is believed to be the first time that all of these technologies have been combined into one process. One of the biggest challenges was to determine the proper dopings to assure low-temperature operation. This required extensive process-level simulations.

7. EXPERIMENTAL RESULTS

Test results for the six basic elements mentioned above are included here. Processed wafers were not received in time to.allow us to report extensive test results in this paper. In particular, we have not attempted to test the main processing array. However, results to date indicate the design approach for the RTP to be valid. The first set of wafers did not have a light shield, principally so that any visible imperfections or errors in the device could be seen during the initial testing phase.

The first device tested was the CCD output amplifier. This was a stand-alone test with an externally-controlled voltage source on the input gate. There was no floating diffusion element which is used in conjunction with the amplifiers on the other test circuits. This test also verified the functionality of the sample/hold feature of the output amplifier. Figures 5 and 6 show the operation of the amplifier. Figure 5 shows the input (top trace) following the output (middle trace). The ground level for the input is the bottom of the scope. The ground level for the output is the third graticle up from the bottom. The bottom trace is the sample/hold waveform which was disabled for this photo. Figure 6 shows the sample/hold enabled. The traces in Figure 6 are the same as those in Figure 5 except the DC offset of the input signal was increased. This shows that the output is indeed sampled and held.

The next device tested was the single latch. This indicated that the input voltage needed for switching was greater than the output high voltage. This was shown to be directly due to a high threshold voltage for the pass gate. By increasing the voltage on the pass gate (above normal TTL), the output high voltage was increased sufficiently to switch the next stage. This allowed data to be passed down a string of such latches. Figure 7 shows a photograph of the latch output. The top trace is the TTL active-low enable (clock) signal. The lower trace is a superposition of the TTL input and the unbuffered output with ground position being the lower level of the input. The photo shows that on the first enable the input is high and the output goes high where it is held until the next enable when the input is low. At this point the output goes low. It remains low because the next enable occurs when the input is low

The basic CCD operation was tested using a four-stage, four-phase CCD structure. These are the same stages used in the output multiplexer of the main device. This CCD empties its charge into a floating diffusion which is connected to the input gate of the output amplifier. Initial indications are that a small input voltage produces an output swing of 400 mV to 500 mV at which point the output amplifier saturates. It is possible that the capacitance on the floating diffusion node is smaller than simulated, restricting the voltage swing to the output amplifier.

The fourth phase of testing concentrated on the clock drivers. These were found to be non-functional. SPICE simulations using the experimentally-determined threshold voltages confirmed that this.was again to due to a higher-than-expected-threshold voltage. Other wafers have been found with acceptable thresholds which will yield functional clock drivers according to SPICE simulations.

The two most complicated structures are the differencer and the comparator. The differencer, although exhibiting the differencing operation, is limited by the low charge handling capacity and low voltage swing of the output amplifier. Its operation in the actual device where it simply empties into another CCD gate would probably be much better.

The comparator test structure contains a layout error in the input structure which prevents the application of a voltage to one of the CCD gates. This in turn prevents the introduction of signal charge into one side of the comparator. We have done some testing assuming that the input signal (as opposed to the threshold signal) is constantly zero and the comparator acts as expected. The layout error does not occur on the main array.

In conclusion, the testing results do not indicate any fatal design or concept errors. The main processing array, though expected to have a low yield, should be capable of full operation. In the least, it may require different clock or DC voltages than what it had been designed for. The design approach for the RTP seems to be validated.

8. PROPOSED ENHANCEMENTS

We are currently designing two enhancements to the baseline processor. The first enhancement is on-chip pixel-by-pixel gain compensation to be used with an indium-antimonide detector array. This would effectively multiply each pixel charge by some predetermined constant and would be performed either before or after the frame differencing operation. The second enhancement is a compact, CCD-based Laplacian window operator. This uses charge regeneration, splitting, and summing to perform the operation.

9. ACKNOWLEDGEMENTS

This chip is being designed, fabricated, and tested under the Real-Time Processing for Staring Receivers contract from Wright Laboratory, Wright-Patterson AFB, Dayton, Ohio. The technical director at Wright Laboratory is Dr. Richard Sanderson.

The design assistance of Dr. Eric Fossum of the Jet Propulsion Laboratory, and Paul Kim (now with Rockwell), Paul Vu, Jeff Pinter, and Dr. Richard Bredthauer of Loral Aeronutronic (formerly Ford Aerospace) is acknowledged. We are using the differencer and comparator circuits developed by Dr. Fossum and his group while he was at Columbia University. Paul Kim provided the 'sea of poly' concept used in the fanout, serpentine memory, and fanin structures. He also provided the CMOS output buffers and the core of the clock drivers. The chip is being fabricated at Loral Aeronutronic (Newport Beach, CA).

The donation of the platinum-silicide array by Hughes Research Center (Carlsbad, CA) is acknowledged.

10. REFERENCES

1. Eric R. Fossum and Richard Clark Barker, "A Linear and Compact Charge-Coupled Charge Packet Differencer/Replicator", IEEE Transactions on Electron Devices, Vol. ED-31, No. 12, pp. 1784-1789, December 1984

2. El-Sayad I. Eid, Charge-Coupled Device Focal Plane Processors, PhD thesis, 1990

3. R. E. Colbeth, N. A. Doudoumopoulos, E-S. Eid, S. E. Kemeny, A. Montalvo, and E. R. Fossum, "High Performance Charge Packet Comparison for Analog CCD Applications", Columbia University, Center for Telecommunications Research Technical Report 095 (1988)

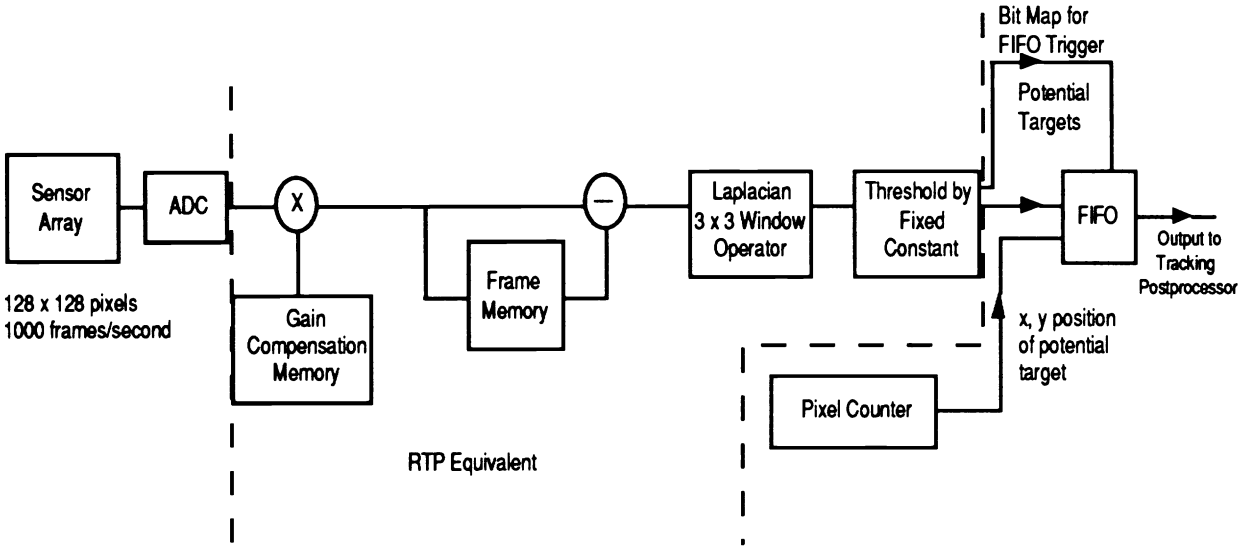

Fig. 1. Implementation of RTP functions with digital processing - data flow/functional diagram.

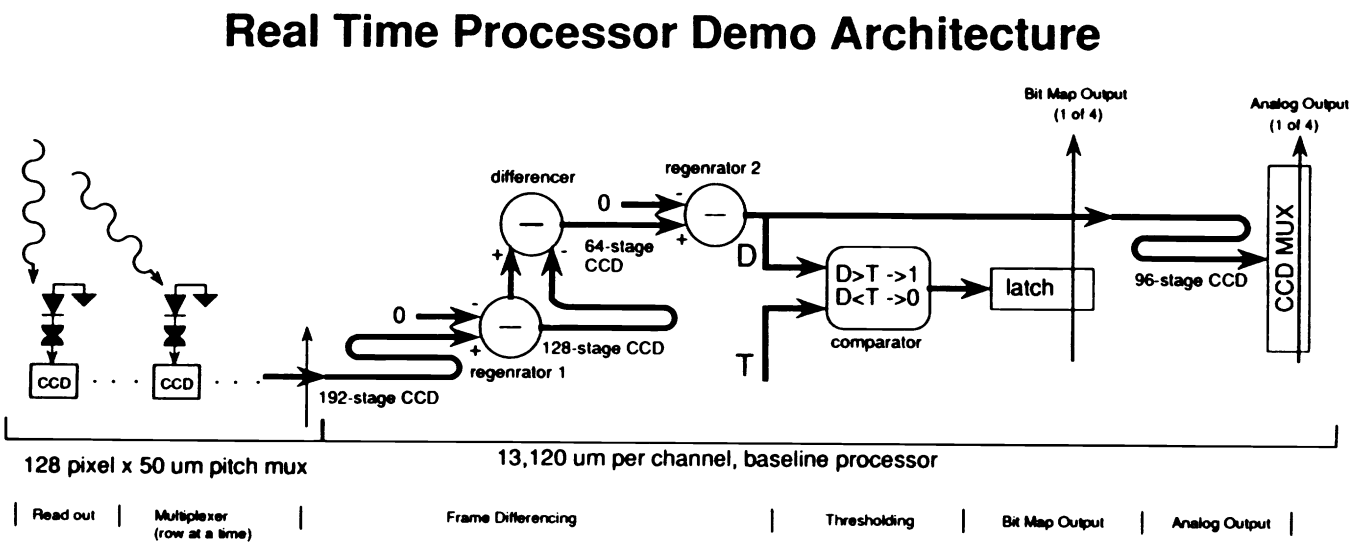

Fig. 2. Block diagram of processor channel.

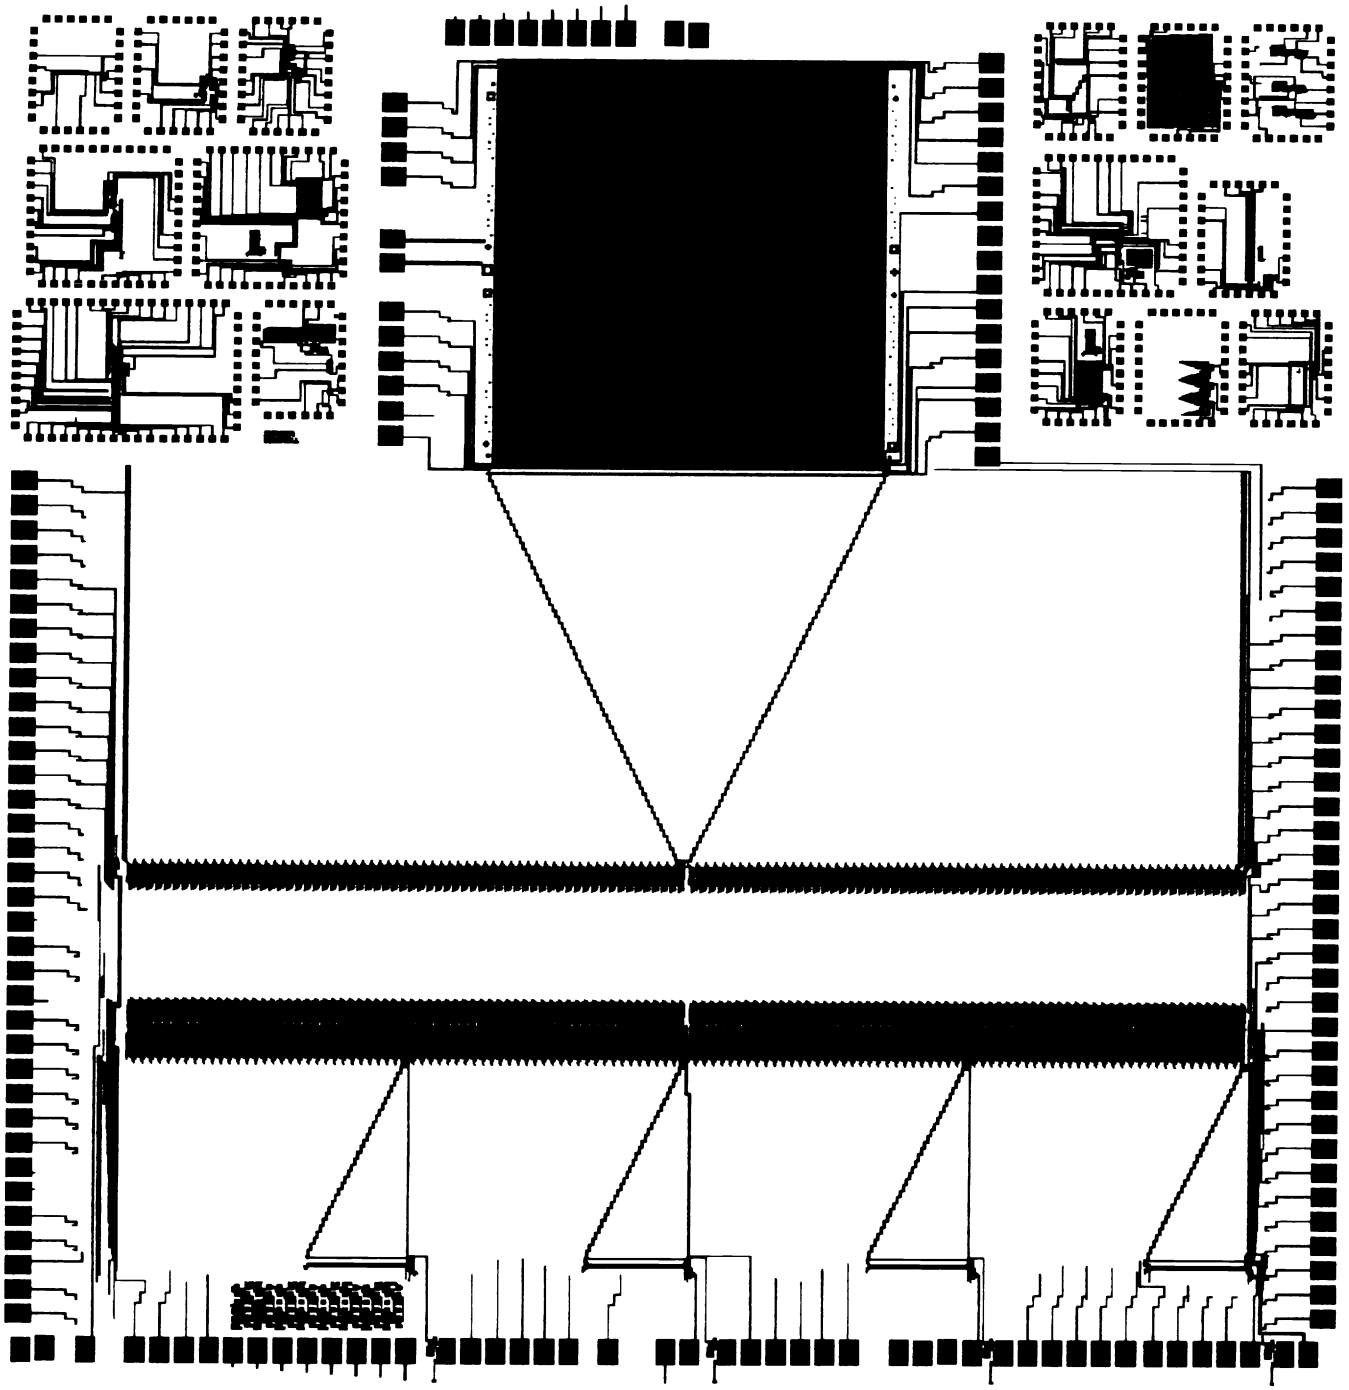

Fig. 3. First metal chip plot showing location of 128x128 multiplexer and pads.

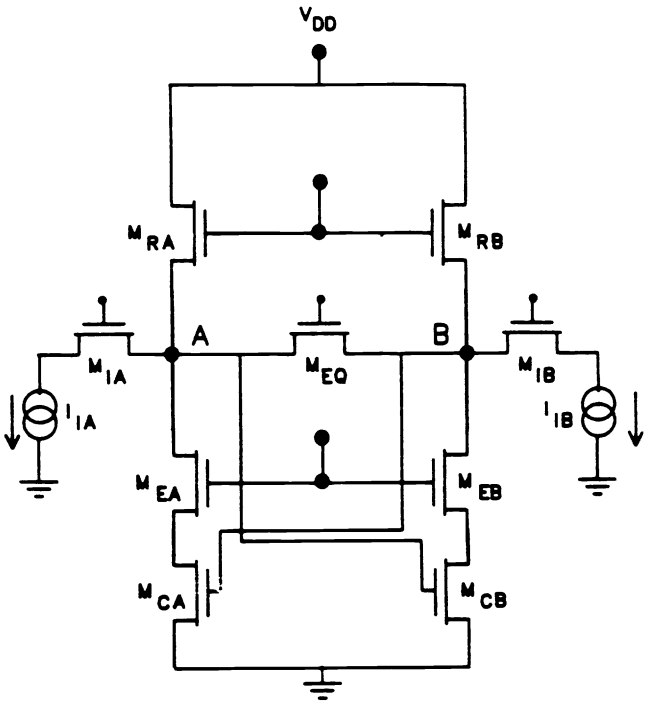

Fig. 4. The comparator schematic as simulated.

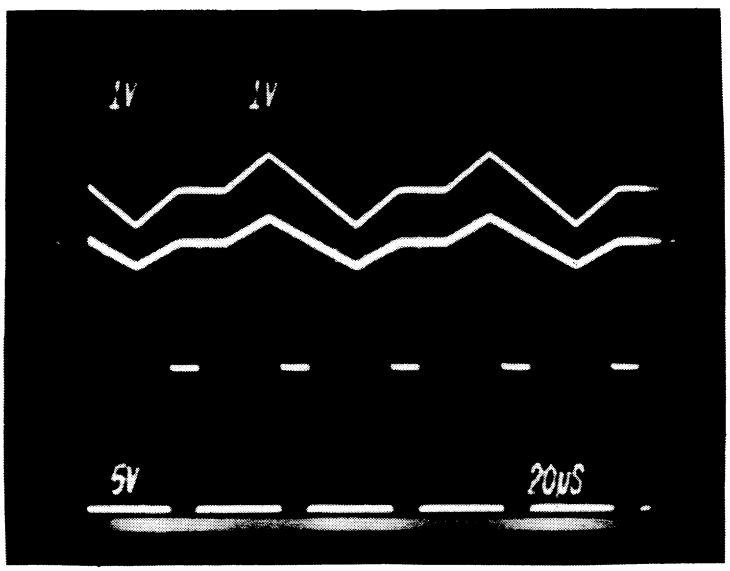

Fig. 5. Output amplifier output.

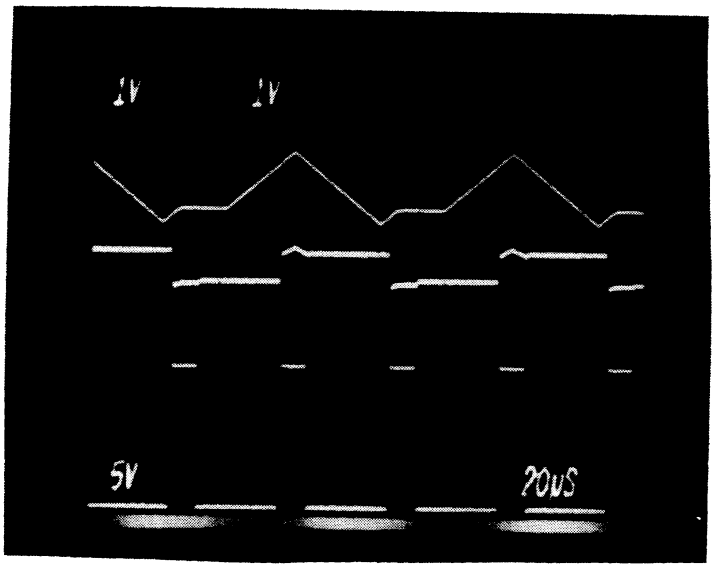

Fig. 6. Sampled output amplifier output.

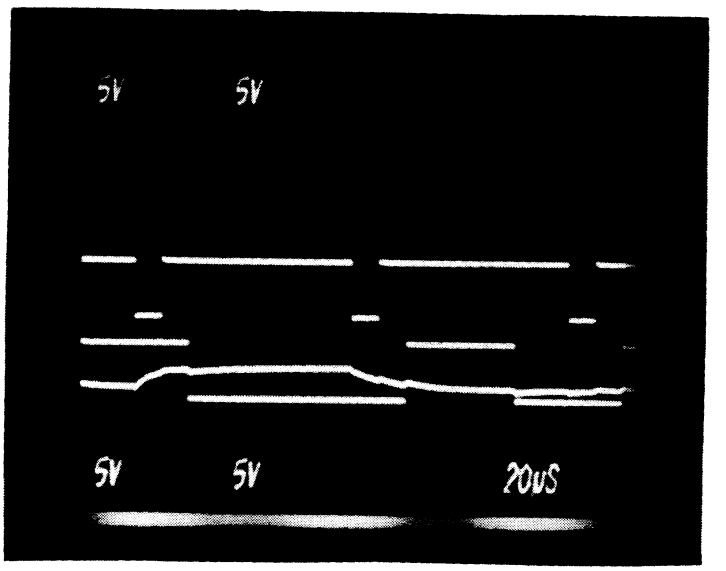

Fig. 7. Single latch output.