Autonomous Data Acquisition using the ACTEL 1020 FPGA

Ron Novick
EDO/Barnes Engineering Division, Shelton, CT 06484

ABSTRACT

The Actel 1020 Field Programmable Gate Array (FPGA) is a low power CMOS device equivalent to 2000 logic gates. It is a one-time programmable device utilizing non-volatile PLICETM anti-fuse programming. The 2-micron, two level metal CMOS process has been tested Rad Hard to 400K Rad total dose Si making it suitable for space applications. In the Barnes Dual Cone Scanner Earth Horizon Sensor, it has been used (with an A/D converter) to autonomously sample radiance data and store it in pseudo-dual port memory. The Actel chip is also used to control bus access to the memory. The Barnes implementation replaces many discrete logic components and frees up processor time by putting the data in memory autonomously. This architecture is applicable to any system requiring low power data acquisition in a minimum amount of board space.

INTRODUCTION

The process of data acquisition involves the representation of a continually varying signal by discrete digitized samples. These samples must be made available to a signal processor; the data often must also be referenced to either time or some other discrete event.

The simplest implementation of data acquisition is to program the processor itself to make use of a clock or event driven interrupt to initiate the data sampling, wait for the conversion to complete, and store the data. A separate reference such as time might be stored with the data to 'tag' it, and/or the location where the data is stored might be used to represent a time or other reference.

This implementation requires only the processor, a Sample and Hold and A/D converter and data memory (RAM). Although this is the most efficient real estate solution, it may represent an unacceptable burden to the processor. The A/D data output time is generally slower than the bus cycle time of the processor and requires the processor to add additional wait states when fetching data. Also, since the A/D converter is generally mapped as an I/O device, the processor must perform an I/O read and then a memory write. For some processors, (particularly the Mil Std 1750), an I/O cycle is significantly slower than a memory cycle. When the processor requires an interrupt to initiate the data sampling, the overhead software to do the interrupt handling must also be executed.

For a 1750 processor operating at 20MHz clock speed, each data sample could take as much as 10 microseconds of processor time for doing nothing but reading and storing one data sample. In contrast, if the data acquisition is performed independent of the processor, the processor requires only 200 nanoseconds to read the data sample from memory. For a 20 KHz sampling rate, at least 20% of the processor's throughput would be spent acquiring data. To read the data directly from memory would only require .4% of the processor's throughput.

If the data could be sampled and stored independent of the processor, the data acquisition processing time could be used more effectively for signal processing or control. A single Actel 1020 chip can be used to perform the data acquisition function and allow the processor to read the data directly from memory in segments rather than one sample at a time. This paper describes a design for the Actel FPGA including all of the timing, logic and register functions needed to perform data acquisition autonomously.

The Barnes Model 1350 Dual Cone Scanner with Sun Fans utilizes this technique. It is a scanning Earth Sensor with two 2.5 degree circular fields of view at different pointing angles relative to the scan axis. It also utilizes a visible light detector with a scanning V-slit field of view to detect the sun and the moon. Both the visible and IR channels are sampled at 4 kHz each,

alternately using the same A/D. A 1750 microprocessor on a VME Space Bus is used to process the signals to detect the angular position of the Earth, moon or sun in each of the channels, with respect to the sensor reference axes. This system is sighted as an example for the application of Autonomous Data Acquisition throughout this paper.

REQUIREMENTS OF AUTONOMOUS DATA ACQUISITION

Sample and Hold

For high accuracy requirements and/or slow A/D converters, the signal waveform must be sampled at an instant in time. That value must be held by the Sample and Hold device until the conversion process is completed. The time instant in which the signal is sampled is often required by the signal processor as well. The Sample and Hold device must be triggered to hold the sample before the A/D conversion is started to allow time for the output to stabilize. The device must then be held in the Hold state until the conversion is completed.

The Actel chip must provide the Sample and Hold strobe and generate the timing to activate the strobe prior to conversion and hold it until conversion is complete.

A/D Conversion

The A/D Converter is the device which takes the DC level of the Sample and Hold device and generates a digital binary value proportional to the magnitude of the held signal. By sampling at a minimum of twice the highest frequency component of the signal, the original signal can be reconstructed from the array of samples according to the Nyquist sample criteria.

The A/D device must be triggered to begin conversion. Typically, it will signal when the conversion is completed and must then be interrogated to output the value before the next conversion is started. A Select line and a Read/Convert line are typically used to command the A/D to begin conversion, output data on its tri-state data lines or remain idle.

DMA Controller

The data must be read from the A/D device and stored in memory. Memory storage can be accomplished either by a microprocessor or a Direct Memory Access (DMA) controller. Both require the ability to generate the correct address on the memory address bus and generate the memory select and memory write strobes and their proper timing. This paper will describe the implementation of a DMA Controller in the Actel 1020.

Local Data Bus Arbitration

If the data storage is accomplished using a DMA device rather than the microprocessor, a system must be incorporated to arbitrate the control of the data and address bus between the DMA device and other devices requiring access to the bus. There are many different types of arbitration schemes; the type chosen must optimize use of the bus between the different masters.

For the Barnes Dual Cone Scanner, the Data Acquisition and 1553 I/O are handled on the same VME Slave board. The Bus Arbiter must arbitrate requests for the data bus between the VME Slave, the 1553 I/O device and the DMA controller (for data acquisition).

The 1553 and data acquisition subsystem only required single word transfers on the bus at periodic intervals and therefore were given priority over the VME Slave, which could monopolize the bus indefinitely. The 1553 device had top priority due to the strict timing requirements necessary to adhere to Mil-Std. 1553. The data acquisition rate was 8 kHz and only required that the data be stored before starting the next conversion. Therefore, the arbitration scheme chosen was fixed priority. The arbiter's job, therefore, is to chose who requesting the bus has highest priority at any given time. Upon completion of every bus cycle, the arbiter must check if a higher priority request came in. If no request is present, the arbiter defaults to the VME Bus so that any request by the VME Bus will not have to wait for the arbiter to grant it the bus (provided a prior request is not in process). The Bus Arbiter requires a Bus Request and a Bus Grant line for each of the 3 Bus Masters.

Sample Reference Tag

The data sample itself is often meaningless without knowledge of where the sample was taken so that the entire waveform can be properly reconstructed. Time or sample number is often the parameter used to tag the data. Using the sample number requires that the sampling occur at fixed time intervals. This then becomes a very simple implementation if each data sample is stored in consecutive memory locations. However, this is not always enough. Often, the data must also be referenced to an external event.

In the case of the Barnes Dual Cone Scanner, the digitized radiance data had to be correlated to the instantaneous angular position of the scanning head. An optical encoder sends an encoded reference pulse every 5 degrees of scan rotation. The pulse is decoded to detect the 0 degree phase reference once per scan revolution. Each 5 degree increment defines a processing segment for the processor to read and process all acquired data during that 5 degrees. The Actel chip provides 2 registers to reference the data to the scan angle. One register holds the last memory address offset for that 5 degree data set and a time reference of that last sample to the encoder pulse. The second register holds the time from the previous 5 degree encoder pulse to the last one. Since the data samples are taken at regular time intervals, each data point can be correlated to a scan phase angle.

ACTEL 1020 IMPLEMENTATION

Clock/Timing Generation

One advantage of using a programmable gate array is the ability to divide a clock by any integer to derive lower clock frequencies. These other clocks are generated as pulses, like the carry output of a counter, and used as clock enables. The clock enable pulses effectively gate the master clock logically while utilizing the edge timing of the master clock. The key to avoiding timing problems in an Actel implementation is using the clock distribution network for all synchronous logic. In any synchronous design, clock skew must be less than the gate delay in order to insure that no component is clocked after its input has already changed state. In the Actel 1020, as well as most programmable gate arrays, propagation delays are often longer than gate delays. The Actel 1020 architecture has one clock distribution network (master clock) with unlimited fanout and a clock balancing feature designed to minimize clock skew.

Wherever possible, timing signals used in synchronous functions should be generated as a clock enable pulse whose width is the same as the clock cycle. Using this clock enable with the master clock ensures that all synchronous circuits will be clocked with a minimum of skew. The timing signals generated for autonomous data acquisition are:

12 MHZ CLOCK
6 MHZ MASTER CLOCK
2 MHZ CLOCK ENABLE
TRACK/HOLD
A/D SAMPLE
CONVERT A/D
READ TIMING PHASE
REFERENCE PULSE

Clock and Clock Enable Signals

The 12MHz Clock Oscillator frequency was initially going to be used as the master clock. However, after simulating the initial design, it was apparent that some of the multi-level logic had delays equal to or greater than the clock period (83 nanoseconds). This would result in unacceptable race conditions so a compromise was made. The 12MHz clock was divided by 2 and rerouted to the master clock to be used for all functions. The 12MHz clock was retained for limited use only where the high clock speed was required for performance.

The logic requiring the 12MHz clock was limited to 3 distinct functions all involved in minimizing bus cycle time. The Bus Arbitration function required fast switching for which 83nS was acceptable but 166nS was not. The VME Slave timing similarly could not

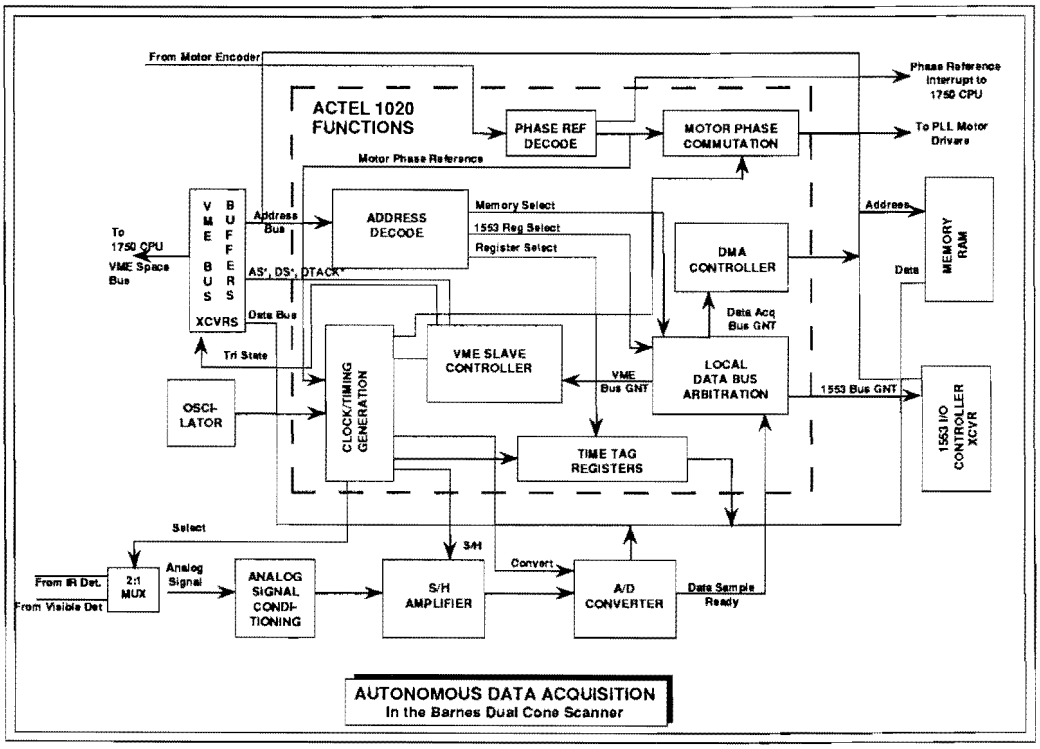

AUTONOMOUS DATA ACQUISITION
In the Barnes Dual Cone Scanner

afford the extra 83nS of slop, both for the local bus access and the system master CPU access time. Then the A/D device required 200nSec for a Data Read, which could be obtained more precisely with $2^1/_2$ 12MHz clock cycles (208nSec) than with $1^1/_2$ - 6MHz clock cycles (250nSec). Each of these functions was limited to 3, 2 and 5 flip flops respectively. By buffering the 12 MHz clock for each function, it was possible to better control the clock skew related to each function. (Clock skew between the functions was not important). Each of these buffered clock nets was declared critical for the automatic place and route and then timing was checked after place and route to check for race conditions.

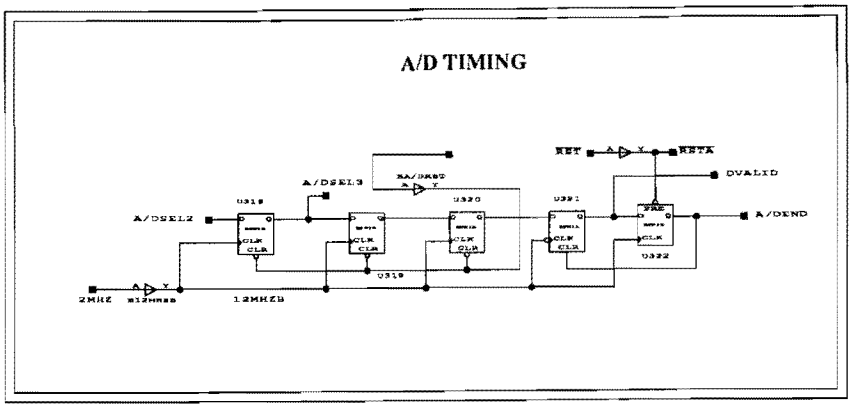

Two MHz was the required timing resolution for the system. Therefore, a 2MHz pulse of one clock cycle duration was generated as a clock enable for the time reference counters. It is also used for other clock divider networks.

Track/Hold Logic

The Track/Hold and A/D Sample Convert logic is shown below.

The 2MHz. Clock Enable signal is divided by 5, 10 and 5 again to produce the 8 kHz sampling rate used in this system. The Track/Hold signal is applied at the same time the A/D converter is commanded to convert. This was acceptable because this A/D device required 100nSec before it began converting, allowing time for the Track and Hold signal to settle. The A/D Convert command is reset after 500nSec, while the Track/Hold signal is held active until the A/D End of Conversion signal is received. The A/D convert is used directly as the A/D chip select (A/DCE).

A/D Read Timing

When the A/D signals End of Conversion (A/DSTS), a request is generated to the Bus Arbiter to grant the bus to the A/D device and initiate a DMA cycle to store the data. Once this is granted by the Bus Arbiter, the A/D Read Timing generates the 200nSec delay necessary for the A/D to output the data (DVALID) and triggers the DMA Controller to write the data to memory. The write strobe timing for the DMA Controller (A/DEND) is also generated from this function so that the same buffered clock is used to guarantee the proper timing.

Phase Reference Pulse

The Phase Reference Pulse is a pulse of 1 clock cycle duration which generates a synchronous clock enable signal from the asynchronous motor encoder signal. This signal is the reference used by the counters for the time tag. Time is measured from this pulse to the time of the last sample and also to the time of the next pulse to correlate the sampled data to the scan phase angle. The motor encoder signal is pulse width modulated and must be decoded to indicate when the 0 degree (once per rev) index reference is crossed. One bit of one of the time tag registers is set when this event is decoded.

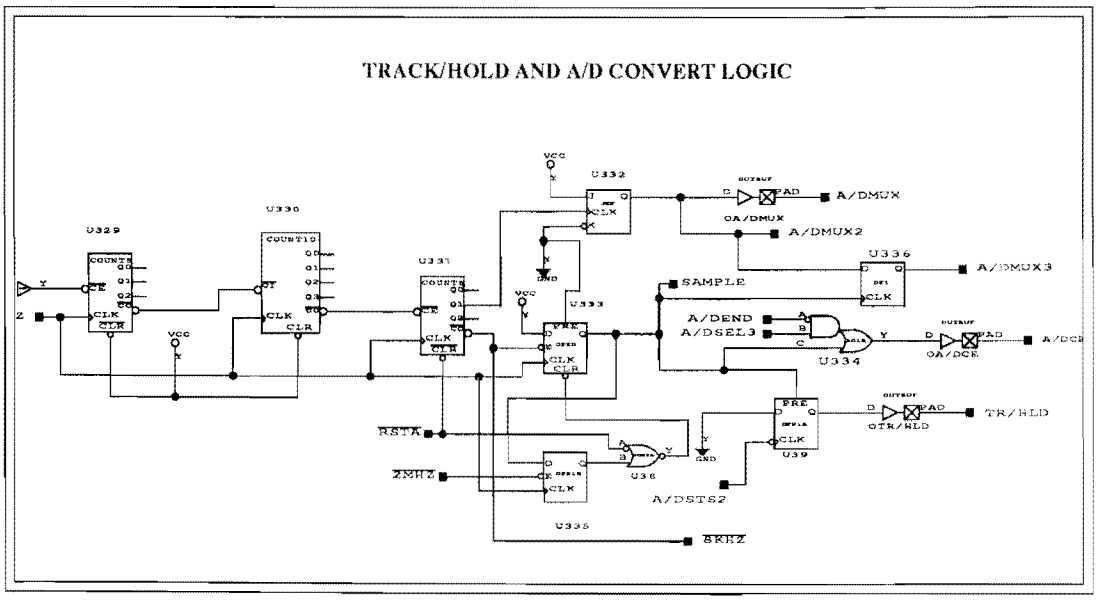

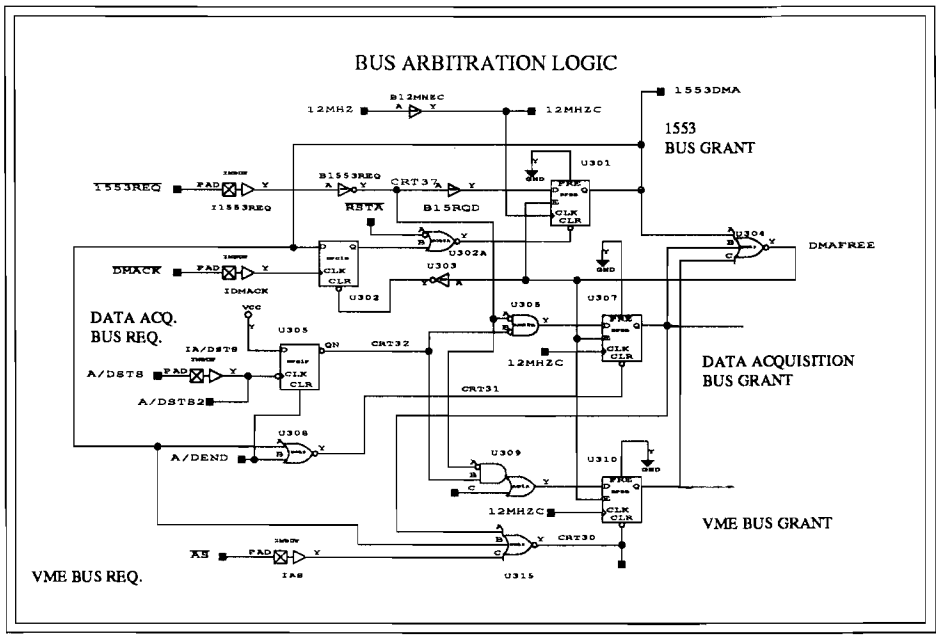

devices to be selected simultaneously, all bus cycles may not begin until one clock cycle after the bus is granted. All higher priority devices, will reset all lower priority devices (flip-flops) when the higher priority is granted the bus (the flip-flop is clocked active). The timing is checked to verify that each device can clear the lower priority devices before the next 12 MHz clock tick in time to prevent it from accessing the bus. Even this fourth technique would not be sufficient alone without the other three. In addition, all output signals are double buffered to help maintain the shortest possible signal routes in the critical arbitration area.

Bus Arbitration

The most difficult function to implement in the Actel chip was the Bus Arbitration. If two or more asynchronous devices request the bus simultaneously, the arbiter must be capable of unambiguously selecting the highest priority requester and denying the others. The difficulty lies in that the combination of clock skew and propagation delay time make it difficult to unambiguously select one and only one device at a time for the asynchronous requests.

The logic circuit shown above accomplishes this by incorporating a variety of techniques. Each of the 3 flip-flops (U301, U307 and U310) represent the bus grant status for the three devices. First, lower priority requests are logically inhibited (at the 'D' input) by higher priority requests. If there was no clock skew and no variation in propagation delays, this alone might be sufficient, but this is not the case. Second, the 12MHz clock is buffered for these three flip-flops and declared critical for place and route to try and achieve minimal clock skew. Third, the outputs of the three flip-flops are or'ed together (U304) and connected to the enable of each flip-flop so that if any one of the flip-flops is active (i.e., a device is granted the bus), they will all be disabled until that flip-flop is reset (i.e., the bus cycle is ended).

Fourth, and most important, acknowledging that it is still possible for two I had to run several iterations of place and route with different criticality assignments to guarantee the timing requirements specified above. This would have been made much easier if the development tools allowed declarations of propagation delay limits. This is a capability being developed currently at Actel called "Timing Driven Placement."

DMA Controller

The DMA Controller circuitry shown below consists of an 8 bit counter for the data acquisition address - tri-state drivers for the address bus - Memory Write, Read, and Enable Strobe Control. The DMA segment offset address is a constant in this design, though it could be made programmable.

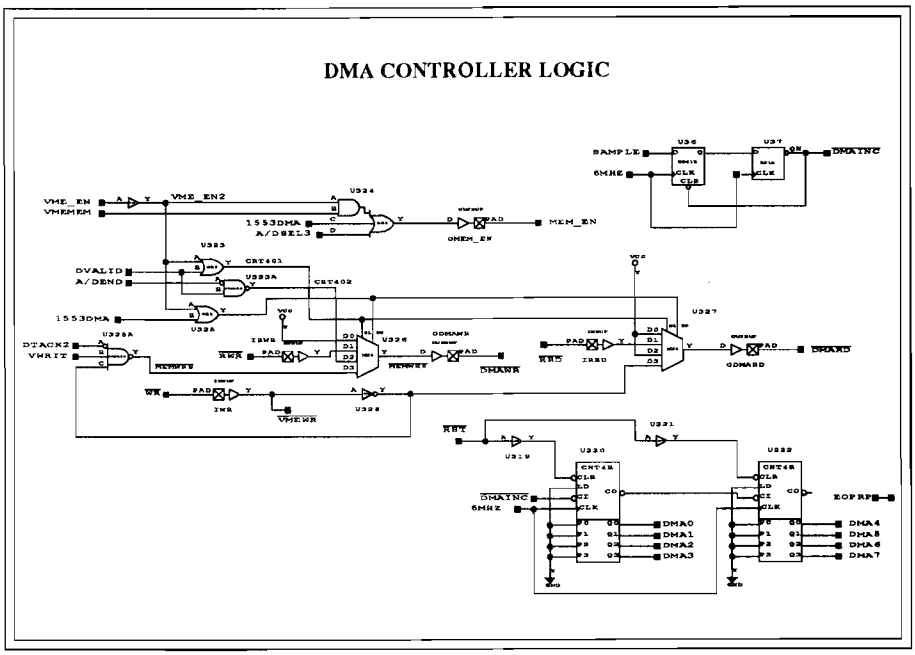

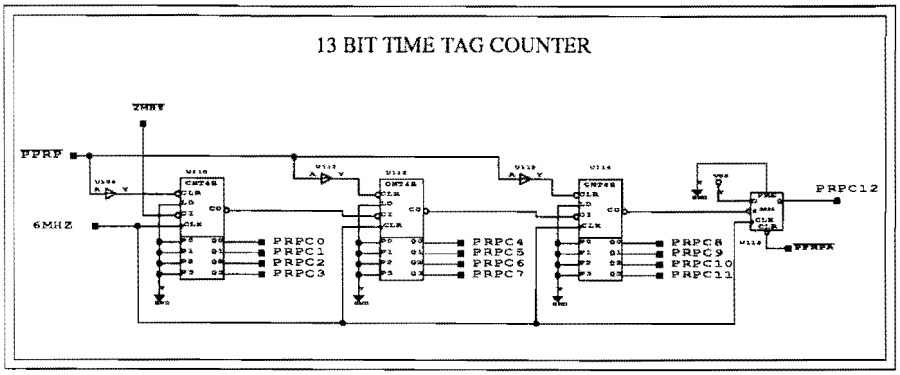

route. Gate delays are also used to help balance the timing; this is generally not a good practice but may be employed when optimum speed must be achieved.

VME Bus memory read times of 185nSec are achievable using 55nSec memory chips assuming no higher priority bus cycles are in progress. This results in a 2 wait state operation for the Barnes 20MHz 1750 processor.

The data is stored in a 256 byte circular buffer starting at the DMA segment offset address. The counter is incremented at every sample conversion using the DMAINC clock enable. The tri-state drivers for the Address bus are easily implemented by the Actel I/O macro.

The Memory Enable strobe (MEM_EN) is generated by a VME memory request, a 1553 memory request or a Data Acquisition memory request. Of these, only the last requires controlling the Address Bus by the DMA Controller. The Memory Enable is kept active for the duration of the Bus Cycle.

The Write and Read strobes are generated for each of these by decoding the appropriate signals from the VME Bus and the 1553 device and generating the proper timing for the strobe. The read strobe timing is not critical and simply needs to be active for the duration of the bus cycle.

The write cycle, however, requires the data be valid on the bus before the strobe is active and removal of the strobe before the end of the bus cycle. The 1553 device generates its own write strobe timing that is simply passed through, but for the VME bus, the VME slave controller generates the write strobe after the data transceivers are enabled (VME_EN2). It terminates it 42 nSec prior to ending the bus cycle. The data has a chance to settle on the bus before strobing it into memory and held active for more than the minimum memory strobe time (35nSec). Timing had to be verified to guarantee that the Memory Enable occurred at least 25nSec prior to the write strobe going active. This was helped by declaring the entire Memory Enable path critical to the place and

Data Correlation (TIME TAG) Registers

A 13 bit counter is implemented as shown above. The counter is reset at every encoder phase reference pulse (5 degrees of scan rotation) and clocked at 2 MHz. The 16 bit register shown below latches the last value of this counter before it is reset.

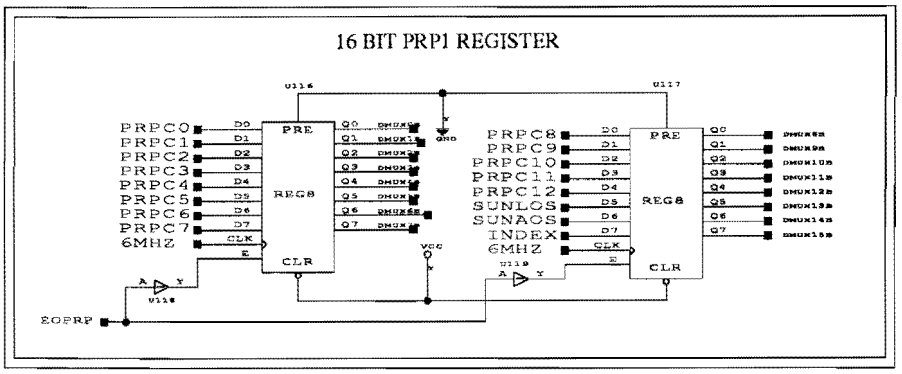

Note the last 3 bits (SUNLOS, SUNAOS and INDEX) are status flags; the INDEX flag marks the 'once-per-rev' encoder mark. The register clock enable is the EOPRP signal. EOPRP is generated 1 clock cycle before the Phase Reference Pulse (~PPRP) which resets the counter.

A second 16 bit register is used to store two values. The first is an 8 bit counter which is reset at every Sample

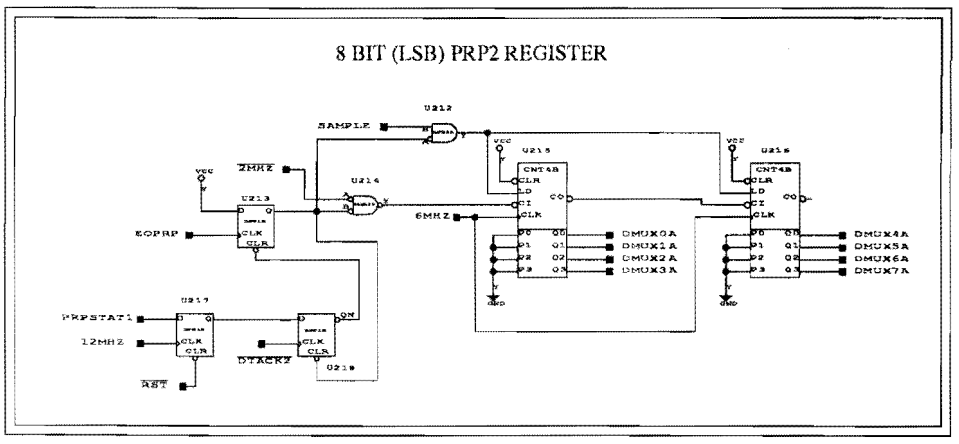

Conversion time and clocked at 2 MHz. The 2 MHz Count Increment (CI) signal is inhibited at the end of the Phase Reference interval (EOPRP) to latch the time value from the last sample to the next Phase Reference Pulse. The counter is restarted after the register is read (PRPSTAT1).

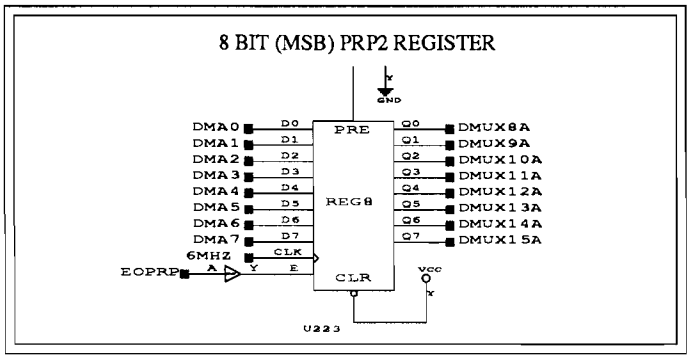

The second value of the second 16 bit register is an 8 bit register which latches the value of the DMA address at the end of each Phase Reference interval (EOPRP).

Two other 16 bit registers are used to time tag other events relative to the Phase Reference Pulse (i.e. the appearance of the sun in the field of view). The 4 registers are multiplexed onto the data bus as shown below:

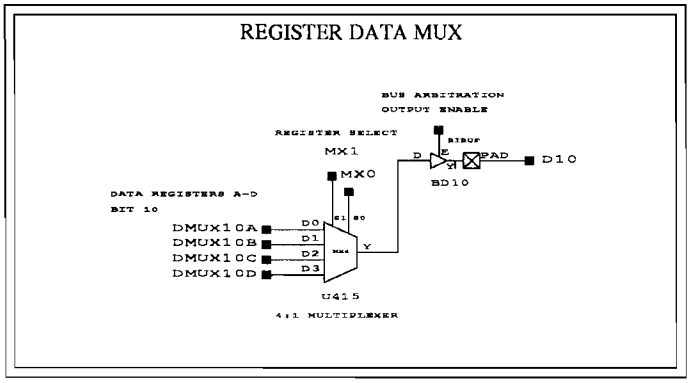

Note that 4:1 multiplexer (for each bit) requires only one logic cell to implement. Each register is addressed by a unique address and is selected and enabled from the address decode circuitry.

Interrupt Generation And Servicing

The Actel 1020 generates an interrupt signal from the Phase Reference Pulse. This occurs at a nominal 3.5 mSec interval during which 28 samples will have been stored. The VME slave module may be used to reset the interrupt according to VME protocol if multiple interrupts are used on the same priority level.

Interrupt processing requires the CPU to read the two Data Correlation registers before the next interrupt (as these will be overwritten). The 28 data samples are stored in a revolving 256 word buffer and need not be read immediately. The data will be overwritten in approximately 9 Phase Reference Pulse intervals (~31 mSec). However, the data is generally read and processed before the next interrupt. The two registers are sufficient to characterize the entire 28 data sample buffer and a memory read is all that is required to upload the data.

By reducing the number of interrupts from 8000/sec to 288/sec interrupt servicing is reduced by 2700%. This frees up at least 20% of the processor's overall throughput.

SYSTEM DEVELOPMENT ADVANTAGES

CPU Throughput

As discussed previously, the time required for the CPU to access the data is reduced significantly allowing more CPU time to be spent on signal processing and control. This is accomplished by delegating the data acquisition task from the CPU to the Actel 1020 FPGA, which autonomously stores the data in memory and generates the appropriate time tags. This reduces the number of service interrupts from one for each sample to one for each 28 samples and allows the processor to access data from memory rather than an I/O device.

Power And Size Requirements

The Actel 1020 is the only chip added to make the data acquisition task autonomous. It is 1.1 inch square for the 84 pin PGA package. The required power was approximately 300mW using the 6MHz master clock and utilizing 90% of the chip's resources.

If the autonomous data acquisition function had to be implemented with discrete digital logic, it would have required 45 to 50 logic chips and taken up a full circuit board by itself.

Concurrent Engineering

By designating use of the Actel 1020 FPGA in the design, only the I/O pin assignments need to be defined initially for the electrical schematic to be drawn. The circuit board layout and design can be done concurrently with the Actel program development. Mechanical/thermal analysis and design can then be done sooner in the development cycle as well.

The reprogrammable nature of the device also allows for last minute changes of the design after the circuit boards are manufactured. Design changes can be made and new chips programmed up to the time the chip is soldered onto the board.

Low Cost Parts And Development

Mil quality parts are available for only a few hundred dollars apiece in small quantities. There is no expensive vendor Non-Recurring Engineering (NRE) associated with ASICs. All the development work may be done on a 386 workstation platform with the Actel Action Logic System and development time is low compared to custom ASICs. The development system comes bundled with Viewlogic schematic capture and Viewsim (or some other) logic simulation may be added to it for complete development capability.

In general, the tools were easy to pick up and work with. Documentation is good overall, though a little light in some areas. The function library is well stocked and documented, and Viewlogic's hierarchical design permitted the creation of function blocks to partition and simplify the design and documentation. Complete placement and routing usually took about 20 minutes on my 33MHz machine. Programming a part after Place and Route takes less than 5 minutes.

VHDL, the Actel Logic Enhancer and Synthesizer (ALES I) and other high level description languages are now available as well.

The only complaint I had about the tools provided was the poor ability to control the routing of critical paths. Paths could only be declared as either not-critical, average, critical, or very critical and the automatic placement and routing would output the best overall combination of what you told it. You could not specify that this path could not carry more than 50nSec of propagation delay which would have been nice. Actel is currently developing this capability called "Timing Driven Placement."

Qualifications of the ACTEL 1020 FPGA

Actel FPGAs use an n-well CMOS technology on 10 micron epitaxial substrate with a 2 micron feature size. The antifuse technology used to program the parts is non-volatile (devices may be programmed only once) and eliminates metalization problems which occur in blowing of standard type fuses for programming.

The Actel FPGA is being processed as a Standard Military Drawing #5962-90964 and is screened to Mil-Std 883C Class B. Radiation testing was performed by the Aerospace Corporation with the following summary:

"Single Event Upsets (SEU) and latchup susceptibilities of CMOS Field Programmable Gate arrays (Actel ACT1010s and ACT1020s) were measured at the Lawrence Berkeley Laboratory 88-inch cyclotron facility with Xe (603MeV), Kr (380 eV), Cu (290 MeV, and Ar (180 MeV) ion beams. The devices did not latchup (at LET > 120, Xe at 60 deg.) and showed a high tolerance to SEU. Moreover, the total dose tolerant level has been reported to exceed 400k Rad (Si). Consequently, these devices may be applicable for use in space." The SEU rate was characterized at about 10-6 upset/bit-day with an LET threshold for upset at 26 MeV-cm2/mg and SEU cross section of 1.0E-4cm2/bit.

The Johns Hopkins University Applied Physics Laboratory also performed radiation testing on the device. Their results showed:

- Functional failure occurring after 500KRad Total Dose
- No latchup occurred at LET=84 MeV-cm2/mg Bromine @60 deg.
- SEU cross section = 2.3E-6 cm2/bit
- LET threshold for upset = 22 MeV-cm2/mg

Future Space Applications of the ACTEL FPGA

Barnes Engineering is already designing the Actel 1020 into its next generation Earth Sensor for small satellites. The new sensor uses all analog signal processing and the Actel FPGA to correlate the analog Earth locator with the digital scan encoder. It outputs a serial digital word corresponding to phase and chord of the Earth signal as well as flag Sun-on-Horizon for each edge. It will also supply a scan speed calibration to correct the Phase and Chord information for any variations in motor scan speed and optionally provide a Radiance correction factor which will allow Phase and Chord correction for radiance variations. The radiance corrections are done in the Spacecraft Processor according to a patented algorithm developed by Barnes Engineering.

The Barnes BX-1750 Space Processor, also developed for its Dual Cone Scanner, utilizes the Bendix BX-1750 single chip microprocessor. The Actel 1020 is used on the processor board to do address decoding and chip selection, VME Bus Master control, and wait state generation. The board achieves a 1.5 MIPS Dais Mix

with floating point throughput at 20MHz. It is Rad Hard to 100K Rad Total Dose Si and is designed with fault tolerant features such as a Watchdog Timer, and EDAC memory (single bit correct and double bit detect Error Detection And Correction). The EDAC function is also performed by the Actel 1020 chip.

Many other Astro-space companies are also looking into applications of this chip for space missions, and I have no doubt it will be seen on many payloads in the near future.

References

"ACT Family Field Programmable Gate Array Databook" by Actel Corp., April, 1990

"Single Event Upset Susceptibilities of Latchup Immune CMOS Field Programmable Gate Arrays" by: R. Koga, D. Baran, K.B. Crawford, S.J. Hansel, B.M. Johnson, D.D. Lau, and S.H. Penzin of the Aerospace Corporation

"Electrical and Radiation Qualification Methods for Field Programmable Gate Arrays in Space Applications" by Prisilla L. McKerracher, Russel P. Cain, Bliss G. Carkhuff, James D. Kinnison of The Johns Hopkins University Applied Physics Lab, presented at the Space Radiation Effects Symposium at Hampton Virgina on March 18, 1991.

"Method of Correcting Errors in Horizon Sensors Caused by Radiance Variations", U.S. Patent #4785169 by Richard A. Gontin, Barnes Engineering Co. Nov. 15, 1988.

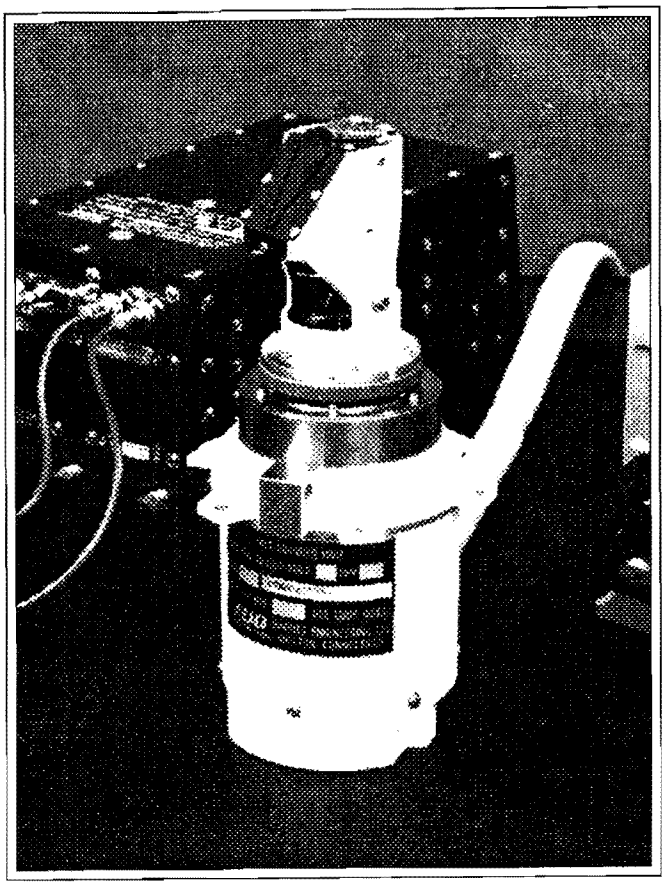

Barnes' Dual Cone Scanner with Sun Fans