Recent Progress in Field Programmable Logic

Alfke, P.

Xilinx Inc., 2100 Logic Drive, San Jose, California 95124
peter.alfke@xilinx.com

I. Beyond Bigger, Faster, Cheaper

Field-programmable logic started out as glue logic between “real ICs.” Over the past decade, however, progress in IC technology has made it possible to implement “real” functions in FPGAs. Now, bigger and faster FPGAs are becoming system platforms that combine several “real” systems functions on a single chip. even microprocessors and memories.

“Bigger” means several million gates, and up to a million bits of RAM, in packages with up to 1156 pins (balls), increasing to 1517 balls in early 2001.

“Faster” means system clock rates of up to 200 MHz and I/O speeds of up to 622 Mbps (800 Mbps and higher in 2001).

“Cheaper” means rapidly decreasing prices. The incremental cost of a Logic Cell (4-input look-up table plus flip-flop) is between 0.4 cent and 3 cents, depending on part type and purchasing volume.

Now that FPGAs can implement complete functions, they must offer not only raw logic in the form of logic cells plus interconnect, but also important system features like on-chip RAM, fast arithmetic (adders and multipliers), sophisticated clock management, and a variety of I/O standards.

The following pages describe features in the presently available Virtex-E devices, and mention some important features of the upcoming (early 2001) Virtex-II devices.

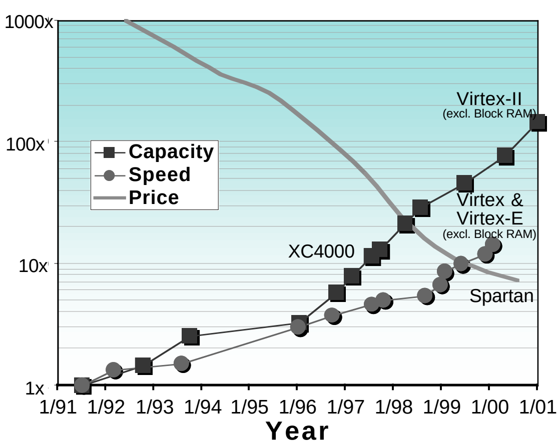

A. On-chip RAM

One FPGA contains up to 65,000 four-input LUTs, each of which can be used as a 16-bit RAM or 16-bit shift register, implementing register files, FIFOs, or dynamically changing look-up tables.

There are also up to 280 larger blocks of dual-ported RAM, each 4096 bits, configurable for different depths and widths. RAM access time is <2 ns, permitting >200 MHz operation, In Virtex-II devices, each BlockRAM will be 18K bits to permit parity-bit storage.

These dual-ported RAMs are ideal for large FIFOs, but can also be used as dual-ported ROMs, implementing state machines, counters (including decimal counters) and code converters. See the Xilinx Application Note XAPP191:

www.xilinx.com/xapp/xapp191.pdf

Using one RAM as a dual Gray-code address generator, a 256-deep FIFO can operate at up to 200 MHz with independent, synchronous or asynchronous, write and read clocks. See the Xilinx Application Note XAPP244:

www.xilinx.com/xapp/xapp244.pdf

The fast I/Os support large external RAMs at a data transfer rate of up to 260 Mbps per pin.

B. Efficient Arithmetic

A dedicated carry structure supports adders, accumulators, and counters with an incremental carry delay of <50 picoseconds per bit. 32-bit circuits can thus run at >150 MHz.

Traditionally, multipliers have been costly and slow in FPGAs, but the upcoming Virtex-II devices provide dozens, or even hundreds, of 18 x 18 combinatorial 2’s-complement multipliers with through-delays of <4 ns (<2 ns for 8 x 8 operation). Since these multipliers are so fast and abundant, they can even be used as efficient barrel shifters.

C. Clock Management

On big chips, clock distribution might easily have become a speed bottleneck, but on-chip digital Delay-Locked Loops (DLLs) solve this problem. They effectively eliminate the on-chip clock distribution delay, and can also be used to eliminate pc-board clock delay.

The clock frequency can be multiplied or divided, generating phase-coherent clock outputs. Slower, phase-aligned clocks can be used to reduce the total clock power budget.

The totally-digital implementation of these DLLs assures robust performance, requires no dedicated power connections or special decoupling, and guarantees <50 ps clock jitter, worst case.

D. Multi-Standard I/O

Xilinx FPGAs come in packages with I/O-counts from 60 to >1000. The proliferation of supply voltages and the increasing emphasis on circuit speed has led to a large number of interface standards. Also, at transition times below 1 ns, interconnect lines as short as 7 cm must be treated as transmission lines that need to be properly terminated, either at the source, at the destination, or both.

Although dedicated level converters and transceivers are available, their use would defeat the main purpose of high-end FPGAs, to reduce pc-board area and maximize performance. For FPGAs with hundreds of signal pins, there is no alternative to direct interfacing.

Virtex-E device pins can be programmed to be compatible with 20 different I/O standards including:

- 3.3V-LVTTL, 2.5V-LVCMOS, and 1.8V-LVCMOS for logic interfaces
- 3.3V SSTL and 2.5V-SSTL to drive series or parallel terminated lines
- 1.5V-HSTL I, III, and IV to drive terminated lines in memory interfaces
- GTL and GTL+ with high sink current and open drain can drive double-terminated lines with 50-Ω pull-ups on both ends
- LVDS and LVPECL differential standards for driving and receiving terminated transmission lines at very high speed

In addition, Virtex devices support double-data-rate interfaces, clocking data on both clock edges.

E. Next Generation

The next-generation Virtex family will use a 0.13µ eight-layer metal copper CMOS process, increase the logic capacity up to 10 million gates, with many 18 x 18 multipliers for high DSP performance, and embedded PowerPC CPUs for distributed processing. System clock rates of >200 MHz and 1 Gbps I/O will be supported.

F. Development Software

New Software offers more than just faster compile time, ease-of-use, and a wider range of alternatives.

New software is also being released to support DSP. FPGAs can achieve superior performance through massive parallelism, but this requires a different design methodology: DSP designers prefer C++ instead of VHDL, and are usually not familiar with FPGA architecture and design flow.

System Generator (to be released September 2000) bridges this gap. It uses MATLAB libraries, and Simulink for modeling and simulation.

Xilinx Blockset (XBS) offers a library of parameterizable DSP functions, visual data flow, system-level abstraction of FPGA circuits, and automatic FPGA code generation.

II. LOW-POWER TECHNIQUES

There are two good reasons for lowering power consumption:

- extend the battery life in battery-powered equipment
- reduce the chip temperature in plug-in-the-wall equipment:

For reliable operation, the maximum junction temperature is 125˚ in plastic packages, and 150˚ in ceramic packages, and performance degrades above 85˚. The best available packages have 10˚C/W thermal resistance without a heatsink or high airflow.

FPGA manufacturers cannot guarantee performance at a specific ambient temperature, because power dissipation (and thus junction temperature) is completely dependent on the user’s design and clock frequency.

A. Low-Power Design Recommendations

In many designs, power is almost evenly divided among clocks, internal logic, and outputs.

- Reduce clock power by minimizing the number of flip-flops driven by a fast clock. Using Clock Enable/Disable does not reduce clock power. Clock gating does help, but can cause hold-time problems.
- Use fast, full-swing input signal transitions to minimize input-buffer current. Avoid floating inputs; one floating input can add 15 mA !
- Control Vcc. Power is proportional to Vcc^2.
- Minimize the number of flip-flop transitions in counters. Gray or Johnson counters are best. Binary counters have twice as many transitions, and LFSR counters have even more.
- Minimize the capacitance of internal nodes by optimizing the design for the highest possible clock frequency. Use aggressive timespecs to force the software to create a tight design with low interconnect capacitance, and thus the lowest power consumption at any clock frequency.

B. Low-Power Design Methodology: A 400-MHz Frequency Counter

This section describes a full-featured, single chip frequency counter that operates at up to 400 MHz, consumes only 130 mW at the maximum input frequency, and occupies 90% of the smallest XC4000 family member, the XC4002XL, or 60% of the newer XCS05XLFPGA device.

The heart of the design is a six-digit decade counter that is driven by a programmable pre-scaler. This pre-scaler is gated by a half-second pulse, and the frequency is determined from the number of input cycles counted in this period.

The time base for the counter is created from a standard 32,768-Hz crystal oscillator. Its output is divided to provide the half-second gating pulse. In a short interval between the gating pulses, the contents of the decade counter are decoded for the 7-segment displays, and the segment states are captures.

The frequency counter has a three-decade auto-ranging capability. At the end of each half-second period, the count value is examined to determine if it is in range. If it is not, the amount of pre-scaling is adjusted for the next half-second period. Hysteresis is built into the auto-ranging circuits to stop any display hunting when the input frequency is at a range boundary.

When the input frequency falls below the auto-range capacity, the display of leading zeros is suppressed. The outputs to the liquid-crystal display are modulated at 128 Hz to provide AC drive directly to the LCD.

A. Semi-synchronous Design

The design uses a cascade of synchronous 2-bit state machines, with each stage clocking the next asynchronously.

Typically, the 2-bit state machine is a modified Johnson counter. The 4-input function generator that precedes each of the flip-flops has three uncommitted inputs that can be used to modify the state sequence.

B. Detailed Design Description

The first stage of the counter is the most critical. At 400 MHz, it is operating at the maximum possible toggle frequency, and the design must, therefore, be kept as s imple as possible.

Consequently, the first stage is an unconditional divide-by-2. The clock-to-setup delay of 2.44 ns permits 400-MHz operation even under worst-case conditions. The flip-flop is located in the leftmost column of CLBs. This location gave the shortest route from the IOB, just 1.1ns.

C. Fixed Divide-by-5 Stage

The residual pre-scaler in the lowest frequency range is divide-by-5. This is in conflict with having the first stage by an unconditional divide-by-2, since five is an odd number.

The solution is a divide-by-2/divide-by-3 counter followed by a toggle flip-flop. This flip-flop is then fed back to control the modulus of the counter, alternating it between divide-by-2 and divide-by-3. The result is that the flip-flop toggles at one-fifth of the input clock with a 2:3 mark-space ratio.

In this case, however, the output is taken directly from the counter. When combined with the first stage, this gives a division ratio that alternates between four and six. This averages to divide-by-5, but with a variable mark-space ratio. Over two-periods, the mark-space ratio is 2:2:2:4. Two clock edges are produced every ten input clocks, and the division ratio is correct.

The count sequence was selected to allow the feedback signal more time to set-up. The control input is "don't care" except at the second clock edge after the toggle-flip-flop is clocked. Thus there are two clock cycles for the feedback path to settle. With a 400-MHz input, 10 ns is available which is more than adequate.

D. Decade Counters

The decade-counter design, is based on the divide-by-5 pre-scaler. The non-binary sequence is not a problem, because LUTs are used as decoders for the 7-segment displays, and any mapping is possible.

E. Results

Using a 3.6-V NiCad battery, the counter operates reliably at 420MHz. As the input frequency varies, the supply current changes from 2 mA with no input to 40 mA at 400 MHz. At idle, the current draw is dominated by the time-base crystal oscillator..

F. Observations

The design is somewhat unconventional in the rate at which its frequency requirements reduce as one moves away from the input. However, it is not that unusual to find small regions of high frequency operation in an otherwise moderate frequency design. The frequency counter demonstrates that, with only a minor amount of manual effort devoted to the high-speed regions, the whole design can easily be implemented in an FPGA.

The complete description of this design may be obtained at:

www.xilinx.com/xcell/xl32/xl32_47.pdf

III. Radiation Characterization, and SEU Mitigation, of Virtex FPGAs

A. Introduction

Field programmable SRAM-based gate arrays (FPGAs) are usually the chosen platform for real-time reconfigurable computing. This technology is driven by the commercial sector, so devices intended for the space environment must be adapted from commercial products.

To evaluate the on-orbit radiation performance expected from FPGAs, total ionizing dose, heavy ion and proton characterizations have been performed on Virtex devices fabricated using epitaxial silicon. The dominant risk is Single Event Upset (SEU), so upset detection and mitigation schemes have also been tested for effectiveness.

This section discusses the radiation performance of Virtex devices, and covers TID, SEL, and SEU. Static and dynamic SEU characterization has been done with both heavy-ion and proton radiation.

B. Technology Considerations

The Virtex FPGA is an SRAM based device that supports a range of configurable gates from 50k to 1M. It is fabricated on thin-epitaxial silicon wafers using a commercial mask set and the Xilinx 0.22µ CMOS process with 5 metal layers.

SEU risks dominate in most applications. In particular, the reprogrammable nature of the device presents a new sensitivity due to the configuration memory. The function of the device is determined when a bitstream is downloaded into the device. Changing the bitstream changes the design's function.

While this provides the benefit of adaptability, it is also an upset risk. A configuration upset may result in a functional upset. User logic can also upset in the same fashion seen in fixed logic devices. These two upset domains are referred to as configuration upsets and user-logic upsets.

Two features of the Virtex architecture help overcome upset problems. Firstly, the configuration bitstream can be read back from the part while it is in operation, allowing continuous monitoring for an upset. Secondly, partial reconfiguration shortens the upset recovery time.

C. Radiation Testing

The space radiation effects of most importance for this work are tolerance to total ionizing dose and single event effects including latch-up and upset. The XQVR300, 300,000-gate Virtex device, was used for testing. Because this technology scales in complexity like SRAMS, it is typical of the entire family.

1). Total Ionizing Dose Tolerance

Total dose testing has demonstrated tolerance in the range of 80 to 100 krads(Si). Testing was done at both high and low dose rates using 60Co sources.

In-situ power supply current measurements were made throughout the course of the radiation exposure. Figure 1 below shows the power supply current monitor traces indicating the onset of TID degradation. Over this range of dose there were no significant changes noted in either AC (timing) or DC parameters, indicating relative stability of the surface MOS thresholds.

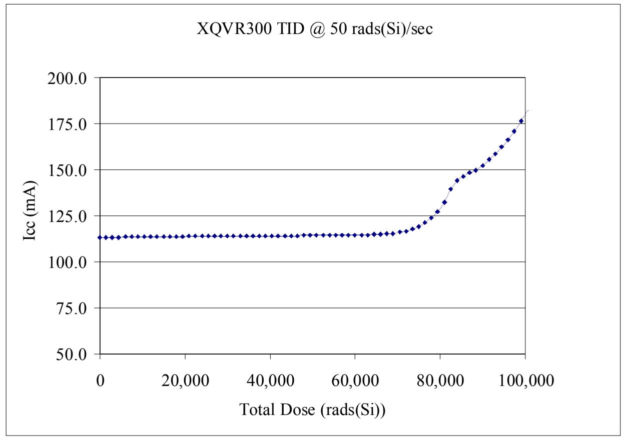

Figure 1: High dose rate performance.

2). Heavy Ion Static SEU & SEL Characterization

Heavy ion characterization was conducted using the cyclotron facility at Texas A&M. Latch-up testing showed immunity to latch-up at an LET of 125MeV-cm^2/mg using gold ions with a fluence of 108 ions/cm^2 indicating no risk of latch up.

Upset testing at the bit level was measured with the resulting cross-section indicated in Figure 2.

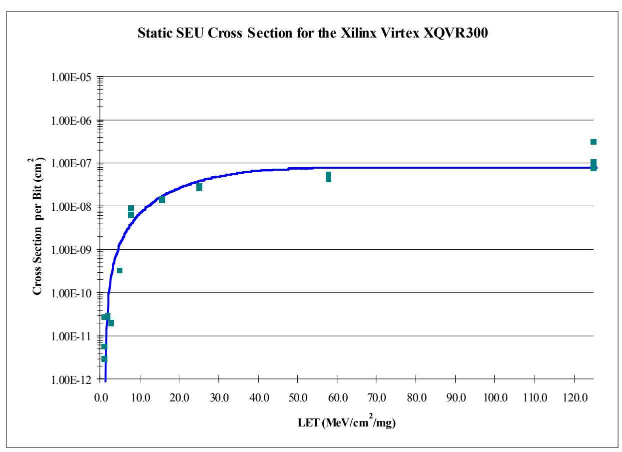

Figure 2: Static heavy ion bit upset cross-section vs. LET

The capability to write and read back the configuration bit stream allowed each routing bit, logic block flip-flop, memory cell, and other storage locations of the device to be individually monitored for static upset sensitivity.

The observed LET threshold was between 8 and 16 MeV-cm^2/mg and only occurred if the fluence exceeded 105 ions/cm^2. Therefore, the device cross-section for this upset mode is very low (<1 E-5 cm^2) relative to the total cross-section of the part and there is a very small probability of occurrence on-orbit.

3). Proton- Induced SEU Testing

Because of the low threshold LET, proton upsets are possible and a similar static bit characterization was performed using the proton beam at UC Davis. The bit cross-section is presented in Figure 3.

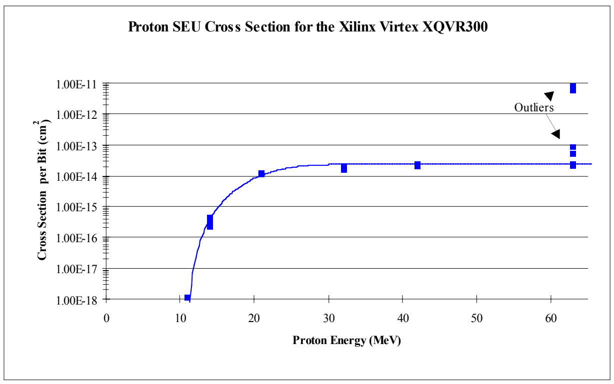

Figure 3: Static proton induced bit upset cross-section vs. proton energy for the Virtex FPGA

4). Discussion of Upset Modes

Upsets in this FPGA can be grouped into three categories: configuration upsets, user logic upsets, and architectural upsets. The physics is the same for all, of course, but the observability and consequences vary.

Configuration upsets occur in the configuration memory and can be detected by readback. The likelihood of failure depends on which bit is upset, and the specific design utilization of the device resources.

Most static bits in the device are accessible via readback. In the case of the XQVR300, there are 1.465M bits in the configuration bits stored, and the cross-section per bit for heavy ions and protons is indicated in Figure 2. Accordingly the static bit cross-section for the part is equal to the product of the number of bits and the cross-section per bit. Of course. the actual cross-section will be less because not every bit upset will be significant in a given design.

The user logic contains elements that are not directly testable for upset through the bitstream. Although most of these elements are accessible through the bitstream, their contents are subject to change due to normal logic operation. These elements include block RAM (BRAM), logic-block flip-flops (CLB-FF), and I/O-block flip-flops (IOB-FF).

Operational upsets can only be mitigated with redundancy in the user's logic design. Observability is limited unless the user design can capture an event. Accordingly, several designs need to be tested to develop useful metrics for these.

Architectural upsets occur in the control elements of the FPGA (e.g. configuration circuits, JTAG TAP controller, reset control, etc). SEUs in these elements are often only detectable indirectly by observing an upset “signature” and associating it with a control element function.

There are two main objectives behind understanding the upset rate and the contribution of these different categories. Firstly, one wants to understand all the possible mechanisms that introduce functional errors. Secondly, to assess the severity of the upset problem, one needs to understand its the frequency and its consequences. These factors determine the cost of mitigation measures and where they are most effectively directed.

D. Mitigation of Single Event Upsets

Two techniques can be used to mitigate SEUs: triple module redundancy and bitstream repair. Triple module redundancy inserts redundant logic into the design to vote out an upset as it occurs in the configuration or even in the user logic.

Bitstream repair uses the fact that configuration readback does not interfere with device operation. Any detected error can be repaired by rewriting the complete configuration, or by using partial reconfiguration.

The paper referenced below describes these methods and their testing in detail

E. Summary & Conclusions

The results of this radiation characterization program show that the Virtex FPGA meets TID and SEL requirements for many orbital applications.

The utility of the device for orbital remote sensing data processing will depend on the mission requirements. The processing performance and survivability of the device are encouraging, but more work is needed to find the source of the dynamic cross-section remaining after mitigation.

For more detailed information, see the paper *Radiation Characterization, and SEU Mitigation, of the Virtex FPGA for Space-Based Reconfigurable Computing* by Fuller, Caffrey, Salazar, Carmichael, Fabula.

This paper may be obtained at:

www.xilinx.com/appnotes/NSREC-2000XPaper.pdf