

THE FUTURE OF FIELD-PROGRAMMABLE GATE ARRAYS

by Peter Alfke, Xilinx, Inc., San Jose, CA (email: peter.alfke@xilinx.com)

Field-programmable logic is ideal for customized digital designs. Like microprocessors and memories, it offers the well-known advantages of very high integration: high complexity and density, small size, low power consumption and cost, and high reliability. On the other hand, programmable logic avoids the problems of ASICs: high Non-Recurring Engineering (NRE) costs, long delays, complex testing issues and the increasingly difficult electrical issues of deep sub-micron ASICs.

1. TYPES OF PROGRAMMABLE LOGIC

Simple Programmable Logic (SPLDs or PALs) were introduced >20 years ago, and are now an insignificant, rapidly shrinking part of the \$2B market for programmable logic.

Complex Programmable Logic (CPLD) devices make up 35% of the market. These devices inherited the AND-OR structure from PALs, but offer more inputs and outputs and better sharing of product terms. Pin-to-pin delays are very short, making CPLDs best suited for wide decoding, synchronous state machines, and counters. The design software is simple, easy to use, and it compiles very fast.

CPLDs are inherently limited in size, and offer relatively few flip-flops. The architecture cannot be expanded to large arrays. CPLDs have a fairly high static power consumption, caused by their wired-OR interconnect structure with many read amplifiers. Only the CoolRunner family (formerly Philips, now Xilinx) offers ultra-low static power consumption.

FPGAs, 53% of the market, have a more ASIC-like architecture with many flip-flops and distributed routing.

A small subgroup of FPGAs uses antifuses to control their interconnect structure. Consequently, these devices maintain their configuration when powered down, they power-on instantly, and they require no external configuration memory. Their internal flip-flops are as sensitive to radiation-induced single-event upsets as any other CMOS storage element, but the logic is fairly immune to radiation problems.

Anti-fuse based FPGAs are one-time-programmable (can be configured, i.e. programmed, only once), and this programming takes several or many minutes. Due to the specialized processing steps, these devices cannot migrate to the newest and most advanced CMOS processes, and they do not offer multi-100,000 or million gate capability, and they probably never will.

Antifuse FPGAs serve a niche market and are only offered by two small manufacturers.

The most successful and fastest-growing programmable device families are the so-called SRAM-based FPGAs. These devices store their configuration (program) in on-chip latches that in turn control pass transistors. Logic tends to be implemented in 4-input look-up-tables (16-bit ROMs). SRAM-based FPGAs offer the highest logic capacity and the highest flip-flop count. The devices can be configured in milliseconds, and may be reconfigured an unlimited number of times. Since they use a standard CMOS logic process, they migrate quickly and easily to the most advanced technology pioneered by the microprocessor industry. The configuration must be reloaded whenever Vcc is being reapplied. This is, however, a major strength of the architecture, since the devices can easily be reconfigured with a new and different program, even after installation.

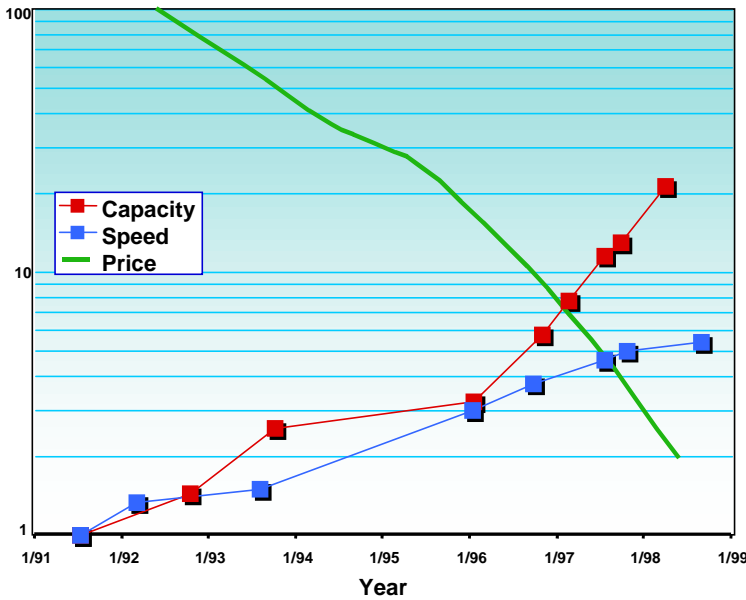
2. SYSTEM DESIGN OPTIONS

- Microprocessors offer greatest flexibility and high functional versatility, but they are too slow for many tasks.
- Gates, MSI, and PALs are inefficient, inflexible and really outdated.
- Dedicated devices and chip sets are powerful and often inexpensive, but offer no design flexibility.
- ASICs (gate arrays and standard cells) offer highest complexity and speed, but suffer from high NRE cost, design effort and risk.
- FPGAs offer flexibility, fast-time-to-market. Dynamic reconfigurability is a unique advantage. Their speed, size, and cost are now approaching those of ASICs.

3. ASIC PROBLEMS

As ASICs are migrating to deep sub-micron technology, they are getting less attractive. NRE cost is driven up by the larger number and increased complexity of their masks. The larger wafer size and smaller die size forces the manufacturer to increase the minimum order quantity. The high-end mainstream ASIC suppliers prefer to deal only with a few, very high-volume users. Low-tech ASICs find themselves in direct competition with advanced FPGAs. In mixed-signal (analog/digital) applications, ASICs have an unchallenged advantage.

FPGA History (XC4000)



- > 20x Bigger
- > 5x Faster
- > 50x Cheaper

Figure 1

Process Technology Evolution

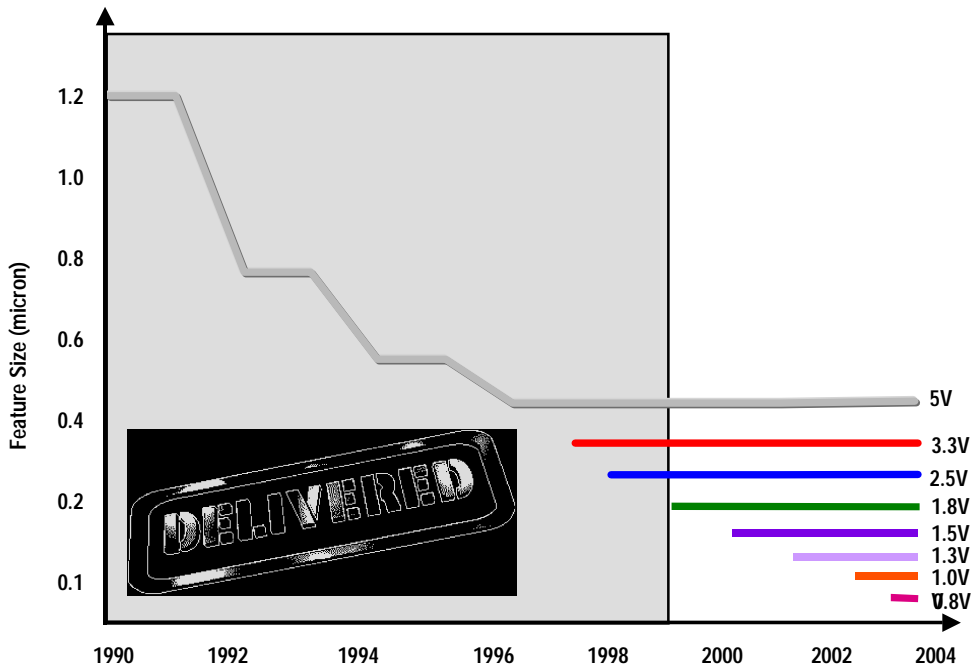


Figure 2

4. FPGA EVOLUTION

The following pages describe the present state and the future of SRAM-based FPGAs. The user community expects a wide choice of device sizes, from 5,000 to several million gates, at speeds up to 100 MHz and above. Design time and effort must be reasonable, and the FPGA supplier must offer and support a wide choice of advanced cores with guaranteed functionality and performance, and must provide powerful synthesis and simulation tools.

Looking back in time, the industry's most successful FPGA family (XC4000) has made tremendous progress between 1991 and 1998 (see figure 1):

The devices got five times faster, the largest available device increased in complexity (gate count) by a factor 20, and for a constant complexity of 10,000 gates, the price dropped by a factor 50.

These historical trends will continue in the future. The coming years will see larger devices, from the present 1 million gates to 2 million gates in late 1999, to 4 million in 2000, 10 million in 2002, an even larger ones in the following years.

The present speed capability can be characterized by >200 MHz on-chip RAM, 200 MHz interface to external RAM, 155 MHz SONET and also 311 MHz bit-serial interfaces, and 66-MHz PCI compliance. The present performance will double by 2002.

5. FPGA PROGRESS

FPGA progress is driven by three independent forces:

- IC technology provides smaller geometries and thus faster transistors and lower cost per function. Better defect density on the wafer makes it possible to manufacture larger chips with acceptable yield.
- FPGA architecture is improved by incorporating system features and by providing a better hierarchical interconnect structure.
- Design methodology is improved with more and better cores, more capable and user-friendly design tools, and faster compile times. The new tools allow a modular, team-based design, and even a distributed design effort via the internet.

5.1. Technology

IC technology has advanced very rapidly during the past 5 years, from 0.5 μ minimum feature size to 0.18 μ today. This offers faster speed and lower cost, but it also means that the 30-year reign of 5V as the only supply voltage is over. Vcc must now be reduced for every new step in the process evolution. Purely by accident, the Volt number is and will be exactly ten times the micron number. (see figure 2)

FPGA technology is essentially identical with microprocessor technology, and thus benefits directly from the fast evolution in that very competitive industry. We use 0.18 μ technology in production today, have 0.15 μ circuits in development, and see a clear road to 0.13 and even 0.10 μ in the future.

Copper interconnect will be introduced in the year 2000, and will be combined with low-k dielectric in 2001, providing lower resistance and lower capacitance for the interconnects, and avoiding metal-migration issues.

FPGA packages have evolved from PLC and PQFP packages with connections confined to the periphery, to ball-grid array packages with increasingly finer pitch. Presently, we offer up to 1156 connections to the chip. The future will see an increasing use of flip-chip packaging technology.

5.2. Architecture

High-end FPGAs must offer more than lots of gates. They must offer a system solution with on-chip memory, a wide choice of interface standards, and must provide sophisticated and robust timing (clock) management.

The Virtex and Virtex-E families offer a 3-level memory hierarchy:

Many (up to 38,000) distributed 16-bit single or dual-port RAMs with sub-nanosecond access time.

Up to 160 versatile 4k-bit dual-port RAM blocks with 3 ns access time and configurable aspect ratio, from 4k x 1 all the way to 256 x 16.

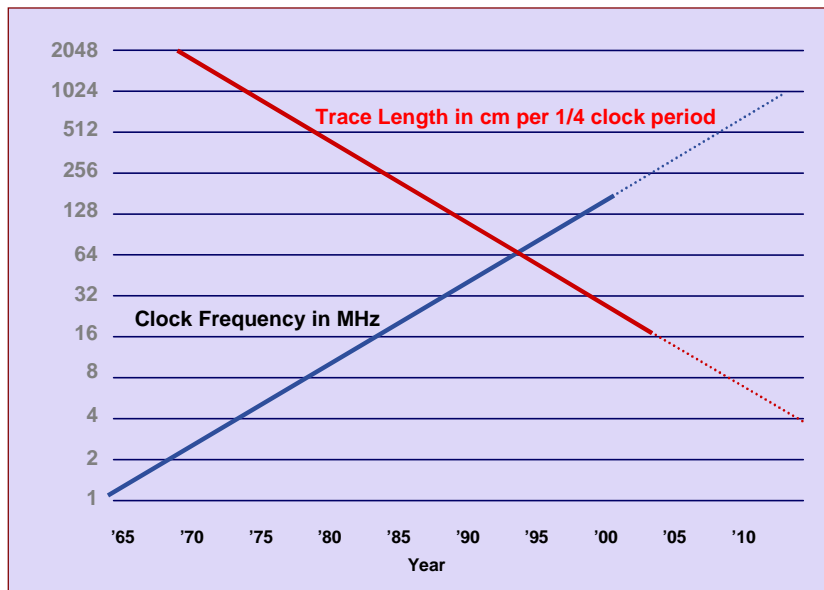
A configurable, fast interface to essentially unlimited external RAM with less than 10 ns access time.

Clock management uses up to eight on-chip digitally-controlled Delay-Locked Loops (DLLs), that can eliminate the on-chip clock distribution delay, de-skew clocks on the board, double or divide the clock frequency, and restore a 50% duty cycle.

As the FPGA implements complete subsystems, it can no longer rely on external level translators. Virtex implements 17 different I/O standards, and the new Virtex-E adds differential LVDS and 3.3 V PECL.

On-chip, Virtex provides a hierarchy of interconnect resources. There are four high-drive, ultra-low-skew global clock nets, each with its own optional DLL, and each capable of driving all flip-flops and registers on the chip. There are 24 additional low-skew global nets for more clocks or other critical nets. All Xilinx FPGAs have bi-directional horizontal Longlines, ideal for on-chip bussing. The many remaining interconnects are segmented, which reduces clock capacitance and thus delay and power consumption.

Moore Meets Einstein



Speed Doubles Every 5 Years...
 ...But the speed of light never changes

Figure 3

Evolution

	1965	1980	1995	2010(?)
Max Clock Rate (MHz)	1	10	100	1000
Min IC Geometries (μ)	-	5	0.5	0.05
# of IC Metal Layers	1	2	3	10
PC Board Trace Width (μ)	2000	500	100	25
# of PC-Board Layers	1-2	2-4	4-8	8-16

- Every 5 years: System speed doubles, IC geometry shrinks 50%
- Every 7-8 years: PC-board minimum trace width shrinks 50%

Figure 4

5.3. Design Methodology

Users demand a more efficient design methodology, driven by high-level languages, and compatible with a variety of industry-standard synthesis and simulation tools. The design effort must be modular, so that several, even geographically dispersed designers can work together on one design.

The internet can be used in several ways. WebFitter allows anybody with internet access to implement a CPLD design on a Xilinx-resident computer. Internet Team Design lets groups of designers share their work over the internet. Internet Reconfigurable Logic (IRL) means that a working FPGA design can be modified, upgraded, tested, or repaired by downloading a new configuration via the internet.

6. RECONFIGURABLE FPGAS

In-system reconfigurability is a unique FPGA advantage with many exciting possibilities. In the design phase, it encourages unlimited experimentation, since mistakes are easily fixed. In production, the system can be customized at the last minute "on the loading dock," or even after it is in operation at its final destination, where the end-user can upgrade a working system. The user can also choose between multiple implementations, and in some cases the system may even reconfigure itself automatically, in a matter of milliseconds, or even microseconds.

Think of an instrument built with FPGAs. Functionality can be changed in milliseconds. One box can serve different purposes at different times. A storage scope can change into a spectrum analyzer, using the same A/D and memory circuits, controlled by a reconfigurable FPGA.

The user can also upgrade or repair the instrument, and thus extend its lifetime, effectively reducing the cost of ownership.

7. CHALLENGES FOR THE USER

Moore's law states that IC complexity doubles every 18 months. A corollary claims that average system speed doubles every 5 years, from 1 MHz in 1965 to >100 MHz in 2000. Unfortunately, the signal propagation speed on a pc-board remains constant at 15 cm/ns. If we postulate that interconnect lines should not waste more than 25% of a clock period, we can calculate a max interconnect length, which shrinks from many meters in the '70s to 30 cm in the year 2000 and 7 cm in the year 2010, when system clock rates exceed 500 MHz. And there is no remedy in sight... (see figure 3)

Higher clock rates demand shorter output rise- and fall-times, about 1 ns today. Interconnect lines longer than 7 cm can no longer be considered lumped capacitive loads, but must be treated as transmission lines, terminated either at the destination or -- if there is only one destination -- at the source. Those 7 cm will change to 4 cm in a few years.

Here is a look at the evolution of digital systems over the 45 year span from 1965 to the future in 2010. It highlights the tremendous progress in the past, but also points at future challenges. (see figure 4)

The rapid increase in the number of metal layers on the IC after 1995 is due to the introduction of Chemical-Mechanical Planarization (CMP) which eliminates the accumulation of surface "bumpiness." Adding a further metal layer now means just a slight increase in wafer cost and a small yield loss.

Power consumption and the resulting rise in chip temperature are a serious concern. Although CMOS consumes practically no static power, the dynamic power is fCV^2 . As the clock rate increases and the chips get bigger, the increasing power consumption is only partly mitigated by a reduction in V_{cc} . Big chips running at high clock rates dissipate >10W and require heat sinks and forced air to keep the junction temperature below 125°C, preferably below 85°.

8. RADIATION TOLERANCE

The new XQR and XQVR series of Xilinx FPGAs avoid latch-up even at 120 MeV cm^2/mg and tolerate >50 krad of total ionizing dose. Single Event Upsets (SEUs) have been investigated and reported, with a primary emphasis on the use in high-altitude flight and Low Earth Orbiting Satellites (LEOS). (See: http://www.xilinx.com/products/hirel_qml.htm#Radiation_Hardened)

SEUs in the configuration latches can be detected by reading back the configuration (which does not interfere with the normal operation of the chip) and comparing it against the original configuration bitstream. SEUs can be corrected by using on-chip triple-redundancy.

9. CONCLUSION

- SRAM-based FPGAs are the fastest-growing segment of the semiconductor industry, sharing technology with microprocessors.
- As standard off-the-shelf components, FPGAs offer fast time-to-market and reduced design effort and risk.
- Density, speed, and cost start to rival ASICs, while avoiding the problems facing the designer of deep-submicron ASICs.
- And finally, only SRAM-based FPGAs can implement reconfigurable systems.

This is the Dawning of the Age of Programmable Logic.