

Nanometers, Gigahertz, and Femtoseconds.

P. Alfke

Xilinx, Inc. San Jose, California
peter.alfke@xilinx.com

ABSTRACT:

This paper describes recent progress in Field Programmable Gate Arrays. The title refers to 90-nanometer manufacturing technology, 11 Gigahertz serial I/O, and a sub-femtosecond capture window causing flip-flop outputs to go metastable.

There are three main sections:

A bird's eye view of FPGA technology, a detailed description of FPGAs in 2004, and two special problems and solutions

I. FPGA TECHNOLOGY

A. Lower cost

Moore's Law is alive, smaller geometries and larger wafers and lower defect density achieve higher yield, and thus lower cost per function. State-of-the-art is 90 nm on 300 mm wafers Spartan-3 uses this technology for lowest cost One LUT + flip-flop did cost \$1.00 in 1990, only \$ 0.002 in 2004.

Rapid price reductions, driven by intense competition

B. More logic and better features:

>100,000 LUTs & flip-flops
>200 BlockRAMs and 18 x 18 multipliers
1156 pins (balls) with >800 GP I/O
50 I/O standards, incl. LVDs
16 low-skew global clock lines
Multiple clock management circuits
On-chip processor(s) and Gbps transceivers

C. Higher speed

Smaller and faster transistors
90 nm technology, using 193 nm u.v. light
Cu interconnect (instead of Al) was easily achieved. Low-K dielectric progress is slow
System speed: up to 500 MHz,
mainly through smart interconnects, clock management, dedicated circuits, flexible I/O.
Integrated transceivers running at >10 Gbps
Speeding up GP logic is getting difficult

D. Better tools

Back-End Place&Route and XST synthesis.
VHDL and Verilog becoming entry point.
IP/Cores speed up design and verification.
Embedded Software Development Tools support architectures and merge HW and SW.
Domain-Specific Languages.
System Generator bridges the gap between Matlab/Simulink and FPGA circuit description.
ASIC-size FPGAs need ASIC-like tools.
ASIC-like size requires ASIC-quality tools.

E. ASICs are losing ground

ASICs are only for extreme designs:
extreme volume, speed, size, low power
Cost of a mask set for different technologies:
250 nm: \$ 100 k
180 nm: \$ 300 k
130 nm: \$ 800 k
90 nm: \$1200 k
65 nm: \$2000 k
plus design, verification and risk

F. 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	12
PC Board Trace Width (μ)	2000	500	100	25
# of PC-Board Layers	1-2	2-4	4-8	10-20

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

G. The ever shrinking circuitry

Number of LUTs + flip-flops + routing that fit on the cross section of a human hair:

- 2000 2 LUTs in Virtex-II (150 nm)
- 2002 3 LUTs in Virtex-IIPro (130 nm)
- 2004 4 LUTs in Virtex-4 (90 nm)
- 2005 8 LUTs = one CLB in 65 nm

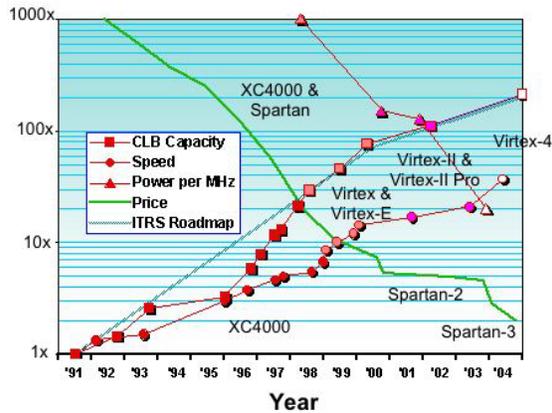
Moore's law is alive and well in FPGAs

G. Middle-of-the-road FPGAs

1990	XC3042	288 LUTs + flip-flops
1994	XC4005	512 LUTs + flip-flops
1998	XC4013XL	1,152 LUTs + flip-flops
2000	XCV300	6,144 LUTs + flip-flops
2002	XC2V1000	10,240 LUTs + flip-flops
2004	XC2VP30	27,382 LUTs + flip-flops
2005	XC4V60	53,248 LUTs + flip-flops

All the same price: One day's engineering salary

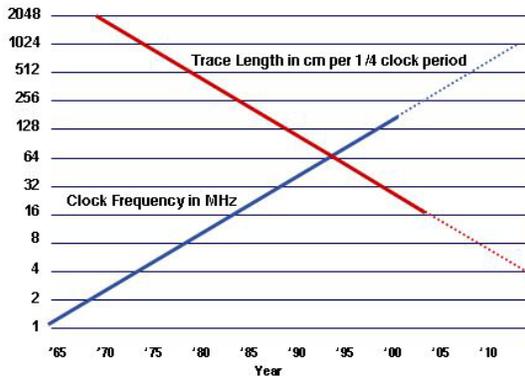
H. Thirteen years of progress



200x more logic in the fabric, 40x faster,
50x lower power per (function x MHz)
500x lower cost per function

Moore meets Einstein

Gordon Moore of Intel predicted that affordable logic density would double every 18 months, doubling speed every 5 years. Albert Einstein postulated that the speed of light is constant.



Speed doubles every 5 years...
...but the speed of light never changes.

I. Higher leakage current

Leakage current = static power consumption
Was microamps, now > 100 mA, even amps (!)
Caused by gate leakage due to 16 Å gate thickness and sub-threshold leakage current
Tyranny of numbers:
10 nA x 100 million transistors = 1 A
Evenly distributed, thus no reliability problem
Sub-100 nm is **not** ideal for portable designs

II. FPGAs IN 2003

1000 to 80,000 LUTs and flip-flops,
Millions of bits in dual-ported RAMs.
Low-skew Global Clocks. DCMs provide frequency synthesis, 50 ps phase control.
18 Kbit BlockRAMs and 18 x 18 multipliers.
1000 to 80,000 LUTs and flip-flops,
millions of bits in dual-ported RAMs.
300+ MHz system clock, 800 MHz I/O,
3+ Gigabit transceivers
Embedded hard and soft microprocessors.
Design security: Triple-DES encryption
VHDL/Verilog entry, synthesis, place and route.
“FPGAs are a compelling alternative to ASICs”

III. VIRTEX-4 IN 2004

90 nm technology, triple-oxide, 1.2-V Vccint
Vcco=1.5, 2.5, or 3.3-V
General-purpose I/O up to 1 Gbps,
0.6 to 11.2 Gigabit/s transceivers
Three sub-families:
V4-LX for logic-intense applications
V4-SX for DSP-intensive applications
V4-FX with PPC and 11 Gbps transceivers
Common architecture for diverse applications

- Higher Performance:
 - 500 MHz for all sub-blocks.
- More Versatility
 - New innovative functions.
- Higher Level of Integration
 - More LUTs, RAMs, multipliers.
- Lower Cost
 - Smaller area = lower cost per function.
- Lower Power per (Function times MHz)
through: triple-oxide gates, multiple thresholds, smaller size, lower Vcc, better design.
Better clocking, less skew, more flexibility.
Better configuration control, better support for partial reconfiguration.
Robust config. cell, SEU tolerant like 130 nm
Flip-chip packaging: lower pin-inductance, stiffer Vcc distribution, less ground bounce.

A. Improved I/O

Supports >50 standards, on-chip termination.
Source-synchronous and system-synchronous.
Serializer/deserializer behind each pin.
Programmable delay available for each pin
> 1Gbps SelectI/O on each pin.
>10 Gbps transceivers (-FX family only).
Source-synchronous I/O improves performance.
Serial I/O saves pins and pc-board area.

B. Faster logic and memory

500+ MHz operation of all internal sub-blocks:
BlockRAM, FIFOs, 32-bit arithmetic, 48-bit
adders and synchronous loadable counters

Up to 72-bit wide memory

4- to 36-bit wide FIFO in each BlockRAM.
Fully independent write and read clocks.
Reliable FULL, EMPTY, ALMOST_FULL,
ALMOST_EMPTY synchronized flag outputs.
FIFOs consume no extra logic and require no
design expertise.

C. Advanced clocking

Proper clocking is extremely important
for performance and reliability.
Most design need many global clock lines
with minimal clock delay and clock skew.
Digital Clock Manager (DCM) provides:
Four-phase outputs,
Frequency multiplication and division
Fine phase adjustment

D. Advanced I/O

>50 Different Output Standards
(strength, voltage, input threshold, etc).
Multiple parallel output transistors
which are either fully on or fully off,
nothing is ever analog, except in LVDS.
Digitally Controlled Impedance =DCI for series-
termination of transmission-line drivers. Adjusts
up/down strength to be = external resistor Two
external pull-up / pull-down resistors per bank.

E. System-synchronous clocking

System-Synchronous when the clock arrives
“simultaneously” at all chips.
Typically used below 200 MHz clock rate.
On-chip clock distribution uses DCM.
Zero clock delay controls set-up time,
and avoids hold time requirements.
“The traditional design methodology”

F. Source-synchronous clocking

Each data bus has its own clock board trace.
Typically used at 200 to 800 MHz clock rate.
On-chip clock-distribution DCM centers the
clock in the data eye, but adds more
unidirectional-only clock lines.
“The only way above 300 MHz”

G. Serial transceiver technology

RocketIO™ Multi-Gigabit Transceiver
8 to 24 per device, 622 Mb/s ... 11.1 Gb/s.
Programmable Features:
64B/66B or 8B/10B Encode/Decode,
comma Detect, Rx and Tx FIFO,
pre-emphasis, receiver equalization,
output swing, on-chip termination,
channel bonding, ac & dc coupling.

IV. VIRTEX-4 COMPARED TO VIRTEX-II

Technology: 90 nm triple-oxide process (three
different oxide thicknesses plus different
transistor thresholds) optimizes the trade-off
between speed, leakage current, and I/O voltage
tolerance . Vccint is now 1.2 V.

Lower static leakage and dynamic current for
any conventional logic implementation, and
drastically lower power when using the new
more highly integrated hard cores (FIFO,
EMAC, DSP slices)

Structure: Radically different, ASMBL chip
layout, arranges functions in vertical columns,
even the I/O. This allows Xilinx to introduce
sub-families with optimized mix between
various functions without upsetting architecture
or software support. The DSP-oriented -SX
family has a much higher ratio of multiply-
accumulators and BlockRAMs relative to the
logic resources in the fabric. Using flip-chip
packages, the ASMBL structure also offers better
power distribution and lower pin inductance.

I/O: Dramatically enhanced capabilities. Each
I/O pin has its own serializer / deserializer
(Parallel/Serial and S/P converter). DDR
interfaces need only one clock and also avoid
any 1-bit latency. A 64-stage individually
programmable delay line in each input can be
used to adjust bit alignment or clock alignment,
ideal for source-synchronous interfaces. BitSlip
supports word alignment. A high-performance
I/O clock can be driven directly from the pc-
board. The larger number (9 to 17 depending on
chip size) of banks gives better I/O granularity,

especially in the larger devices. Configuration now has its own bank.

CLBs: faster, but no significant structural change. To reduce chip area and to enhance performance, only 50% of the slices have LUTRAM and SRL16.

BlockRAM: optional pipeline output registers double performance to 500 MHz. Two BlockRAMs can directly be concatenated for 36K x 1 operation, or for 512 x 72 bit operation with built-in Hamming error correction. The optional hard-coded FIFO controller inside the BlockRAM runs at up to 500 MHz reliably, even with asynchronous read and write clocks. Data bus width, fall-through operation, and the levels of the ALMOST_EMPTY and ALMOST_FULL flags are programmable. The FIFO takes up no fabric area, is much faster, much easier to design, and more reliable than previous soft cores.

DSPSlices: Significantly enhanced from the traditional 18 x 18 two's complement multipliers. Now include a 48-bit accumulator with three adder inputs and very efficient cascade capability.

Clock Management Resources

DCMs are faster, more precise, and have better jitter performance. Two modes: highest speed and finest granularity, or longest delay range. Dynamic reconfiguration of phase shift values and M, D. (But M and D changes still require a DCM reset). Phase-Matched Clock Drivers (PMCDs) provide multiple (divided) clock outputs that are delay-matched.

Global Clocks are routed differentially, thus faster and less sensitive to crosstalk. Duty cycle is better maintained. More capable BUFGMUX. clocking architecture adds high-performance I/O clocks, and a large number of regional clocks.

Available in Virtex-4 FX only:
PowerPC runs 12% faster and has new APU co-processor interface with direct connection to processor instruction pipeline. A third-party floating point unit by QinetiQ provides up to 250 MegaFLOP performance. Enhanced OCM interface includes I/O handshaking.

RocketIO from 0.6 to 11.1 Gbps (wider range than Virtex-IIProX) with additional advanced input equalization options.

EMAC is a new function, integrates the digital portion of 10/100/1000 Mbps Ethernet Media Access Controller (MAC). Saves over 2000 slices of a soft-core alternative.

V. VIRTEX-4 CAPABILITIES

Any type of design runs at >400 MHz.
Pipelining provides extra performance "for free."
Synchronous is best, but 32 clock are available.
Gigabit serial saves pins and board area.
On-chip termination for board signal integrity.
I/O features support double-data rate operation and source-synchronous design.
Popular functions are hard-wired for low cost, higher performance, and ease-of-use:
Microprocessors, clock management, FIFOs, I/O serializer/deserializer, Ethernet MAC etc.

Many pre-tested soft cores are available, some are free, some for a fee.
One-hot state machines are preferred, but MicroBlaze or PicoBlaze may be better.
Massive parallelism enhances DSP.
Up to 1024 fast two's complement multipliers per chip, faster than dedicated DSP chips, but need system-rethinking

VI. CHALLENGES

Technology moves rapidly: 130, 90, 65 nm, multiple Vcc, lower voltage - higher current. Lower Vcc makes decoupling very critical. Moore's law becomes more difficult to sustain. Leakage current has increased significantly, but triple-oxide transistors, threshold control, and clever design provide relief. Signal integrity on pc-boards is crucial. "Homebrew" prototyping would waste money and time. Use evaluation boards instead

VII. BOARD-LEVEL PRODUCTS

AFX basic evaluation boards (chip features)
Low-Cost **ML40x** (~ \$ 500) (system behavior)
ML46x- Memory Eval. Board (interfaces)
ChipScope Pro for real-time debug.
Debugging usually dominates the design effort, needs access to chip-internal nodes and busses. It is often practically impossible to dedicate the necessary extra pins and routing.
Do not waste time "debugging the debugger".
ChipScope Pro has internal virtual test headers, small cores act as internal logic state analyzers. ChipScope Pro provides full visibility at speed. Read-out via JTAG, no extra pins needed.
ChipScope Pro is the best tool for logic debug.

640 MHz Clock Generator.

Direct Digital Synthesis, PicoBlaze for control.
Frequency synthesis generates <350 ps jitter.
External PLL reduces jitter to 100 ps.
1 Hz steps, 1 ppm frequency accuracy.
1000 frequencies stored in EEPROM.
Small, low cost, easy single-knob control.
Next generation will offer 5 GHz in 2005.

VIII. TWO PROBLEMS AND SOLUTIONS

A. Single-Event Upsets in Virtex-II

SEU = random soft error,
directly / indirectly caused by cosmic radiation.
Known problem at high altitude and space,
traditionally not a problem at sea level.
Many tests, papers, show ways to mitigate:
Readback, scrubbing, triple redundancy.
Aerospace designs tolerate the cost/size penalty.

Traditional Test Methods: Vastly accelerated testing procedures, bombarding an operating FPGA at Los Alamos and Sandia Labs. Many SEUs are detected and reported, but there is no agreed-upon conversion factor to “normal” terrestrial operation.

B. Xilinx large-scale test

4 boards with 100 XC2V6000s each,
running 24 hrs/day, internet-monitored
readback and error logging 24 times/day
Locations:

San Jose, (at sea level)
Albuquerque, NM (1500 m elevation)
White Mountain, CA (4000 m)
Mauna Kea, Hawaii, (4000 m)

What's the Real MTBF ?

Measured mean time between SEUs in
XC2V6000 at sea level is 18 to 23 years.
But >90% of config. cells are always unused.
The real Mean Time Between Functional Failure
therefore is 180 to 230 years for XC2V6000
or ~1000 years MTBFF for XC2V1000

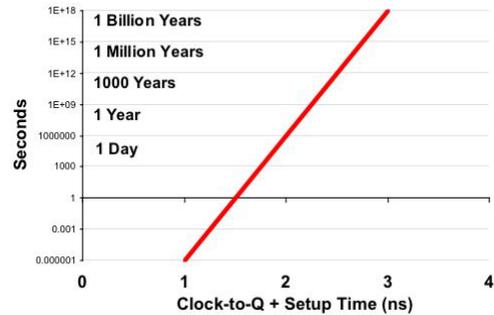
C. Metastability in Virtex-IIPro

Violating set-up time can cause unknown delay.
Potential problem for all asynchronous circuits.
Problem is statistical and cannot be “solved”.
Xilinx published tests in 1988, 1996, and 2001.
Modern CMOS flip-flops recover extremely fast.
Metastability is now irrelevant in many cases.
The small metastability capture window was

tested on Virtex-IIPro devices:

0.07 nanosec. for a 1 ns delay clk-to-Q + set-up
0.07 femtoseconds for a 1.5 ns delay,
etc.

One million times smaller (shorter) for each
additional 0.5 ns of acceptable delay. This
parameter is independent of clock and data rates.



Mean-Time-Between-Failure as a function of
tolerable delay (300 MHz clock, ~50 MHz data).

IX. Conclusion

FPGAs have become cheaper, faster, bigger,
more versatile, and easier to use. They are now
the obvious first choice for the system designer.