20 Years of FPGA Evolution

from Glue Logic
to Major System Component

Peter Alfke, Xilinx
HOT CHIPS 19
August 19, 2007

XILINX®

Before FPGAs: 1965 to 1985

From Gordon Moore to Ross Freeman

"Moore's Law":

"Transistor count doubles every 2 years"

from 32 transistors in 1965 to 1 billion in 2007,
25 doublings in 42 years, and still going on…

An amazingly accurate prediction,
self-fulfilling prophecy, industry-roadmap,
although not a law of physics.

For 42 years, our engine of progress

Evolution: SSI to MSI to LSI

- From Small-Scale Integration, Medium-Scale Integration to Large-Scale Integration, evolving from Bipolar to CMOS
- **Good:**
  - Higher functionality, reliability, lower cost and power consumption
  - Ideal for memories and application-specific ASSPs
- **But:**
  - Higher integration means much narrower specialization
  - Bad for versatility and diversification of the end-product
- Coming to the rescue:
  Microprocessors and **Programmable Logic**

The advantage of LSI, but completely versatile

Makimoto's Wave

1957 to '67 Standard discrete devices, transistors, diodes
1968 to '77 Custom LSI for calculators, radio, TV
1978 to '87 Standard microprocessors, custom software
1988 to '97 Custom logic in ASICs
1998 to '07 Standard Field-Programmable devices
After 2008 ? but Moore's Law will continue…

– *Tsugio Makimoto, formerly Hitachi,*
Chairman of the Technology Board of Sony Semiconductor Network Co.

Field-Programmable Logic Evolution

- **PALs**, user-programmable MSI, (MMI)
  - slow, high power and expensive, but successful
  - 1981 “The Soul of a New Machine” (DG “Eagle”)
- **CPLDs**: re-programmable CMOS (Altera)
  - LSI, lower power, lower price per function
- **FPGAs**, configured by antifuse (Actel)
  - More logic, mux-based, one-time programmable,
- **FPGAs**, configured using latches (Xilinx)
  - “SRAM” –based, volatile, but re-programmable, and:

…uniquely able to take advantage of Moore’s Law

Build a Versatile Machine; Let the User Define the Application

One mass-produced design, many widely different applications

- Johannes Gutenberg, Movable-type printing, 1450
- Joseph Maria Jacquard, Punched-card loom, 1801
- Herman Hollerith, Punched-card tabulator, 1890
- Mauchly & Eckert, ENIAC, 1946
- Ole Kirk Christiansen, LEGO Bricks, 1949
- Intel, Microprocessor 1971
- Apple and IBM, Personal Computer, 1978
- **Ross Freeman, Xilinx FPGA, 1985**

…tremendous flexibility, wide market acceptance

Xilinx FPGAs 1985 to 1999

Ross Freeman's **revolutionary** assumption:

"Transistors are cheap, and getting ever cheaper
Six transistors to emulate a short circuit is reasonable"

Homogeneous array of Look-Up Tables and flip-flops

plus dedicated bidirectional I/O.

Connectivity and LUT content controlled by latches.

Volatile = infinitely reconfigurable, even partially reconfigurable.

Reconfiguration was a new, strange, unproven concept.

Fear of volatility required extensive missionary work

Original family: XC2064 and XC2018

- **Logic fabric:**
- 64 or 128 identical blocks of Look-Up-Tables plus flip-flops.
- **Speed:**
  max 50 MHz clock rate (typically lower)
  - Identical I/O capabilities on all user pins.
  - Surface-mount packages.
- **Perception:**
  - Small, slow, expensive and "different"
  - Risky, from a small, unproven supplier
    "How do you spell Xilinx, how do you pronounce it?"
- MMI (later AMD) signed up as "second source".

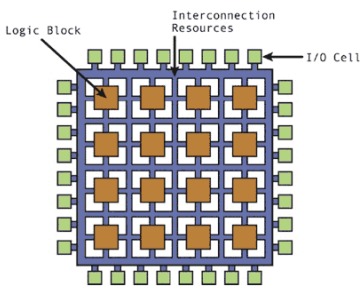

Moderate success after intensive missionary work.

XC3000 and XC4000 families (1987 to 1999):

- **XC3000**: Improved fabric: 4-input LUT,
  - very successful,
    despite rudimentary software support;
    severe routing limitations, and slow speed

- **XC4000**: Added system features:
  - Dedicated carry logic to make arithmetic faster and cheaper
  - LUTs used as distributed RAM, even dual-ported
  - Several sub-families

- **Software** support becomes important
  - NeoCad tries to be vendor-independent, fails, and joins Xilinx

Popular circuits to support uPs and memory

Common Configuration Methodology

After power-up:

- Internal sequential house-cleaning
  - (clearing the configuration store of up to 80 Mbits)
- Then loading (push = Slave, or pull = Master)
  - a serial bitstream or byte-stream,
    coming from external storage, e.g. Flash or µP
- Outputs are 3-stated during configuration
- DONE output signals complete configuration

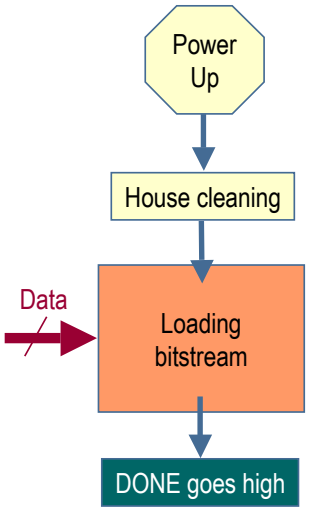

Partial reconfiguration in the newer families

Virtex Families: 1999 through 2007

- Process evolution = higher speed, lower cost.
- **Better architecture** brings big benefits:
  - Dual-ported BlockRAM, first 4Kb, then 18K and 36K
  - 18x25 bit multiplier, 48-bit accumulator
  - Digital Clock Management and PLLs
  - IDELAY, ODELAY with 75 ps granularity
  - Multi-Gigabit Transceivers (3.2 to 6.4 Gbps)
  - Hard Cores: PowerPC, Tri-mode Ethernet, PCI Express
  - System monitor, Analog-to-Digital converter for Vcc and temp

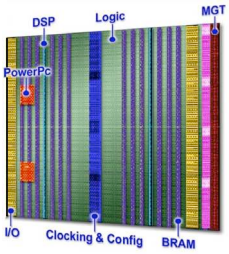

Breakthrough to System-on-a Chip

Designing with FPGAs

- Take advantage of massive parallelism (different from µP)
  - Avoid long chains of combinatorial logic
  - Increase throughput by extensive pipelining
- Look for specific dedicated functions (hard cores)
- Use "left-over" functions creatively:
  - PowerPC or BlockRAM = state machines,
  - Multiplier = barrel shifter, 48-bit accumulator = fast counter
- Appreciate clocking features and I/O versatility, on-chip termination, SerDes, GigabitTransceivers

Design with the FPGA architecture, not against it

FPGA Advantages

Custom design without NRE, no risk, and no manufacturing delay
- offers unlimited design diversification, avoids commodity price wars
- Fully characterized chips, (like µPs,) far superior to ASICs

Fast parallel operation, instead of slow serial
- much faster than microprocessors and DSPs
- abundant flip-flops can enhance speed through pipelining

Offers opportunity for **unlimited modifications**, even in the field
- just like µP code, far superior to all "frozen" alternatives

Can combine a **multitude of diverse functions** in one chip
- logic, arithmetic, cores, clocking, I/O features, multi-gigabit transceivers

Can interface to a **wide variety** of electrical standards
- unmatched by any alternative solution

FPGA Disadvantages

Higher power consumption, slower speed, and less capacity
- than possible with ASICs and ASSPs

Higher cost per chip than ASICs
- but no NRE, thus **lower total cost** for most applications

Higher cost per chip than ASSP
- but the opportunity for design diversification avoids commodity price wars

Requires you to adapt the logic design to the FPGA features
- not as granular as ASICs, but all features are pre-tested and characterized

Software tools are not as advanced as for ASICs
- but are much cheaper

General trend: Away from ASICs, towards FPGAs

One Size Doesn't Fit All

- 3 or 4 sub-families offer functional diversity ('LX, 'LXT, 'SXT, 'FXT)
  - 5 sizes cover 30k to 330k LUTs, a one-to-ten range
  - 3 speed grades x 3 temp ranges x 2 package options
- Two different **business** models: Virtex and Spartan
  - **Virtex:** High performance, feature-rich, large size
  - **Spartan:** Low cost / low power, optimized for consumer
- New technology introduced every 2 years
  - "because it can be done", often ahead of user need

Seventeen FPGAs Devices Delivered on Four Platforms

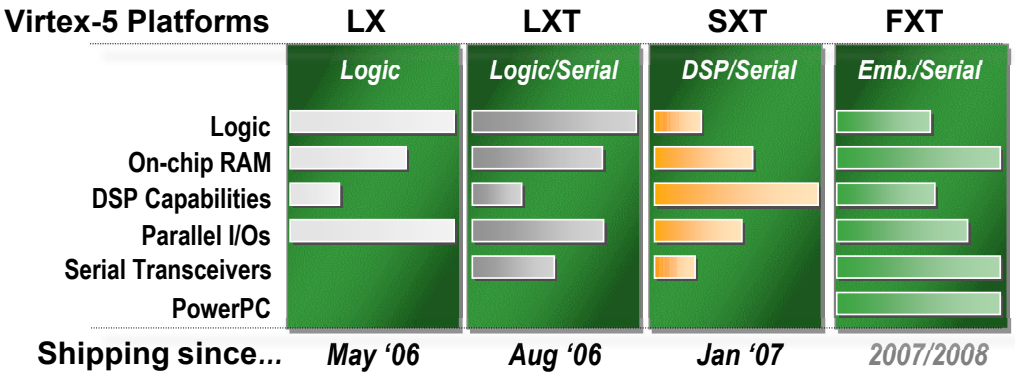

- High-Performance Platform FPGAs
- Built-in PCI Express and Gigabit Ethernet
- Embedded BlockRAMs, DSPs, PowerPC, and Serial Transceivers

Spartan, Spartan-II, Spartan-3 Families

- Emphasis on reduced cost, simpler packages
  - Optimized ratio of logic to I/O for specific applications
  - Growing family, non-volatile and DSP-oriented members
- System features:
  - Subset of Virtex features, optimized for lower cost and power
  - Cost-sensitive market demands narrowly optimized feature set

Going after the consumer market

Spartan Devices Optimized for High-Volume Applications

Addresses specific cost and flexibility requirements

Broaden the Market, 1990 - 2005

200 x capacity, 40 x faster, 500 x cheaper, 50 x lower power

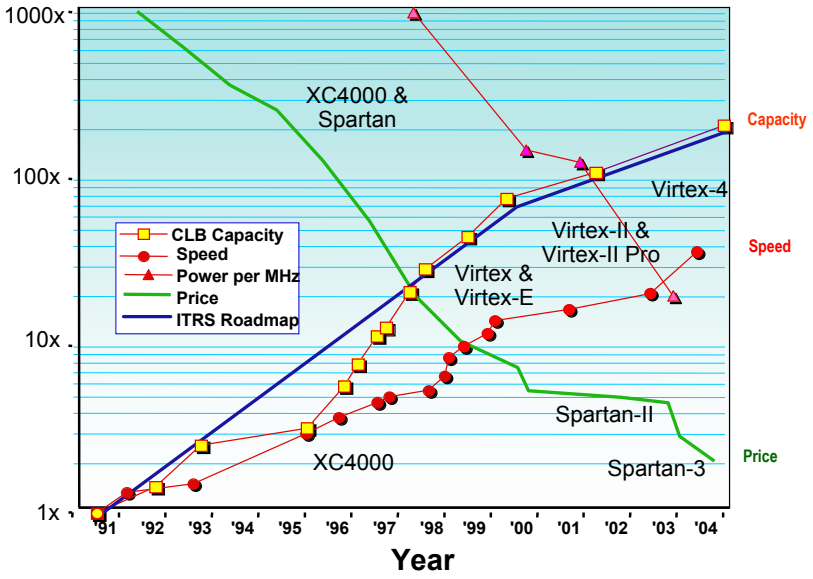

FPGAs are Everywhere…

- For most jobs, the newer chips have sufficient routing,
  - Flip-flops and most blocks have sufficient speed
  - Interconnect delay is the performance limitation.
- FPGAs are always used for their speed (otherwise use µP)
- Users demand powerful, bug-free, easy-to-use tools
  - with large number of available (hard and soft) cores.
- New users are increasingly hardware-ignorant, (HDL-savvy)
  - Synthesizers cannot cope with combinatorial feedbacks, or with set-up time violations, or with asynchronous clocks
  - Multi-clock designs are non-trivial
  - Never use decoders as clocks

Some users attempt designs that are way over their head

"Like Duct Tape or WD-40"

"FPGAs can perform **packet switching**,
they can accelerate **digital signal processing**,
they can **bridge incompatible communications protocols**,
they can play host to entire **embedded computing systems**,
they can drive **backplanes**,
they can act as **reconfigurable computing** elements.
People are building devices that consist of little more than one or two FPGAs plopped down onto a circuit board with a few generic interfaces.
Look at our new reconfigurable **supercomputer**, they say.
Observe our amazing automotive **telematics integrator**, they boast.
See our fantastic **ASIC emulator**, they cry.
We, however, always just see a board with a few FPGAs.
The situation reminds one of the blind men with the elephant
No, more accurately, it reminds one of Duct Tape."
Quoted from:
Duct Tape, WD-40, FPGAs
The Universal Survival Kit
*by **Kevin Morris**, FPGA and Structured ASIC Journal*

God's Gift for Prototyping

Almost all digital breadboarding and prototyping is now done with FPGAs

The fastest way from idea to working model
The easiest way to experiment and get repeatable results
The best way to prove a concept or try out a change
Much faster, cheaper and more convincing than any simulation
The best way to demonstrate audio and video designs,
where a computer simulation may be meaningless
The easiest way to build multiple copies
The best way to give the software folks some hardware to play with

Evaluation boards can avoid costly effort and delays
Re-usable for university educational projects, from freshman to Ph.D. level

We support prototyping, although not as our main business

Xilinx University Board 2005

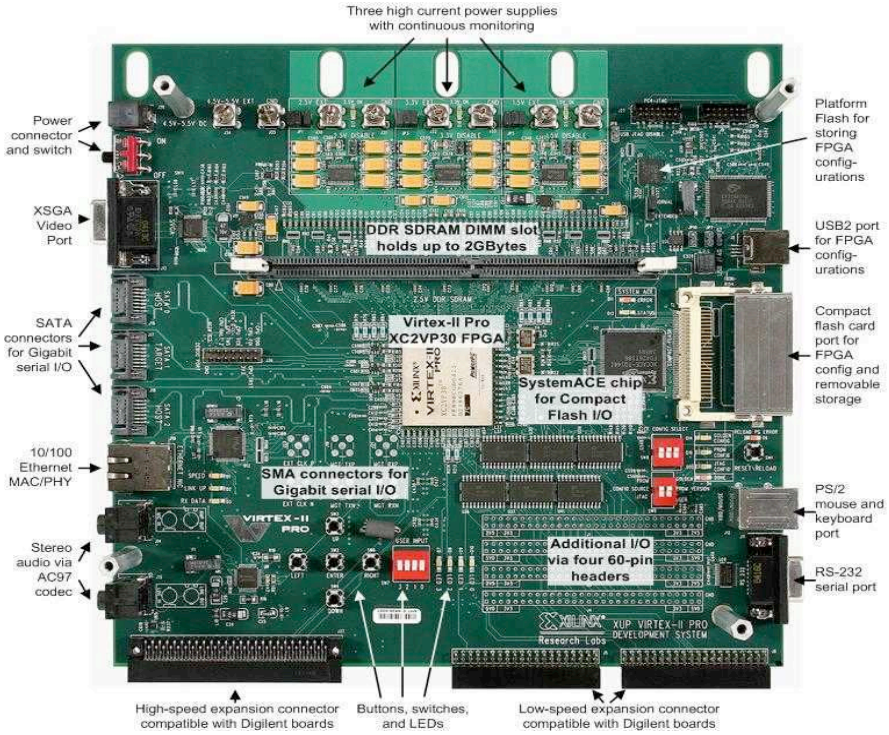

Virtex-5 Evaluation Board 2006

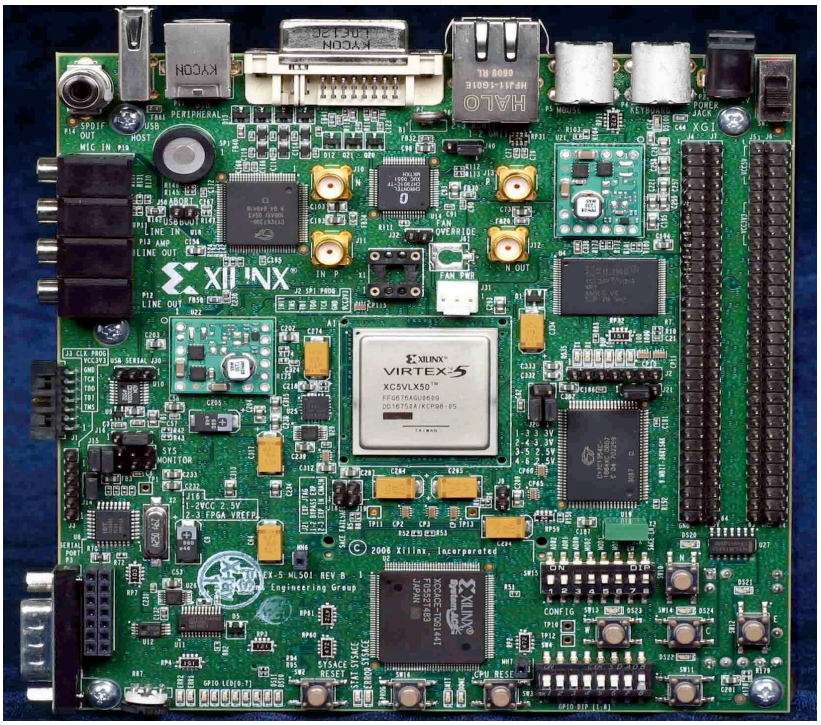

Business Aspects

Second sourcing:

- Vital in the '80s, now irrelevant except for commodities
- Xilinx "Dual-Fab" strategy assures availability

Vendor-Independent Place-&-Route Tools:

- NeoCad could not make any money, joined Xilinx

Patent Battle:

- Xilinx and Altera sued each other
- Only the lawyers got rich

Intense Competition between Xilinx and Altera

- Leads to rapid introduction of better FPGAs

Innovation may be too rapid for some users…

Progr. Logic Market Share

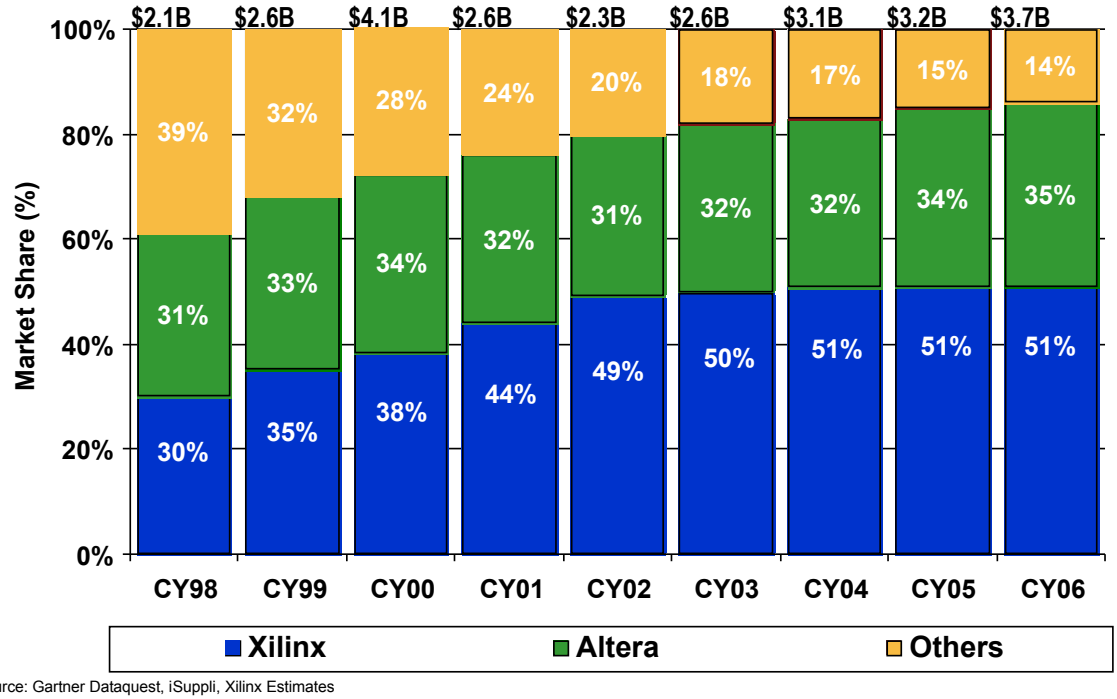

Source: Gartner Dataquest, iSuppli, Xilinx Estimates

A Maturing Industry…

- The two leaders have a combined market share of 86%
  - Xilinx 51%, Altera 35%, leaving 14% for Lattice, Actel, Quicklogic, Atmel…
- All **broad-based suppliers** have given up:
  - Intel, T.I., AMD, ATT, Philips, Cypress, NSC, Motorola
  - "FPGAs demand too much management attention"
- Most **start-ups** vanished:
  - Dynachip, PlusLogic, Triscend, Siliconspice (absorbed)
  - Chameleon, Quicksilver, Morphics, Adaptive Silicon (failed)
- Fab-less might make it easy, but there are strong barriers to entry:
  - Tool familiarity, design re-use, risk avoidance by the users,
  - Economy of scale, patents, technical support, availability of sales channels

Where is the next successful FPGA start-up?

Mainstream Features

- Versatility, high performance, and low cost
- User-friendly and capable tools
- Many available cores, helpful tech support
- Easy (partial) re-programmability
- Signal integrity on the pc-board
- Compatible I/O levels and standards
- Many size, speed, temp, package options

FPGAs replace ASICs and ASSPs

Low-Priority Niches

- Single-chip, non-volatile, instant-on
  - is left to CPLDs, usually for small systems
- One-Time Programmable (OTP)
  - antifuse technology, limited appeal for aerospace
- Ultra-low power for battery operation
  - high performance causes high leakage current
- Limited perceived security
  - Antifuses & Flash are not good enough for serious protection
  - (Virtex devices offer encryption with battery-backed-up keys)

…of marginal interest to the main players

Fab-Less is a Winner…

- FPGAs need leading-edge technology
  to overcome inherent inefficiencies
  compared to ASICs and ASSPs
- Leading-edge fabs are expensive ($2B to $3B)
  - and become obsolete in a few years
- Fabs like TSMC and UMC are very profitable
  - pioneering aggressive technology, spreading the risk
- FPGA companies do R&D, design, test, and marketing
- Leave manufacturing to (mainly Asian) fab partners

Fab-less gives stability, fast boom-time growth, and easy survival in a recession

FPGAs Often Beat ASICs

- Faster pace / shorter product lifetime
  favors field-programmable solutions
- ASICs only for extreme applications:
  - Very high volume, very high speed or very low current.
- FPGAs are often:
  fast enough, big enough and cheap enough.
- FPGAs offer the absolute fastest DSP implementation.
- Dynamic reconfiguration becomes practical.

Practical /affordable FPGAs cause system-rethinking

FPGAs Drive the Process

New process technology drives down costs for all customers.
FPGAs are able to take advantage of new technology faster than ASICs and ASSPs

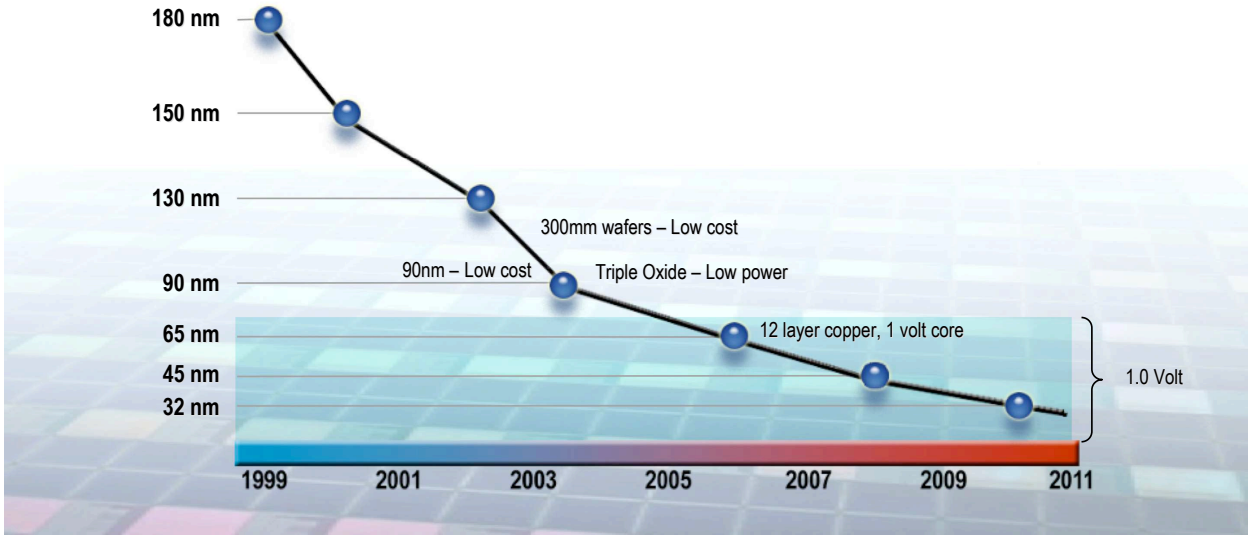

As technology progresses down the process curve, only companies that can amortize their investment over many customers and designs will be able to afford to design their own IC's

Technical Obsolescence

ICs do **not** wear out in normal use
They are available for **many years,**
but they become **non-competitive** in a few years:

> ***"One year in the life of an IC***
> ***equals 15 years in the life of a human"***

2 years old: IC is tops ***(capable, competitive, available)***
4 years old: IC is a senior citizen ***(with certain limitations…)***
6 years old: IC is to be "fondly remembered…"

Modern FPGAs are a great substitute for obsolete devices

Challenges

- Non-homogeneous architecture choices require trade-offs
  - Efficiency, density, speed and low power against versatility
- Distributing and protecting IP cores owned by multiple vendors
  - Ease-of-use while honoring the rights of the IP originator
- Users increasingly ignore hardware / architecture trade-offs,
  - Many rely blindly on smart synthesis packages
  - How to keep a hundred thousand users happy inspite of themselves
- Single-Event Upsets (SEUs) from radiation is a perceived problem,
  - but exhaustive tests (Xilinx "Rosetta" project) **quantify** the issue.
  - **65 nm devices are actually more robust than older devices**

(Static) power consumption is the major concern

CMOS Power Trends

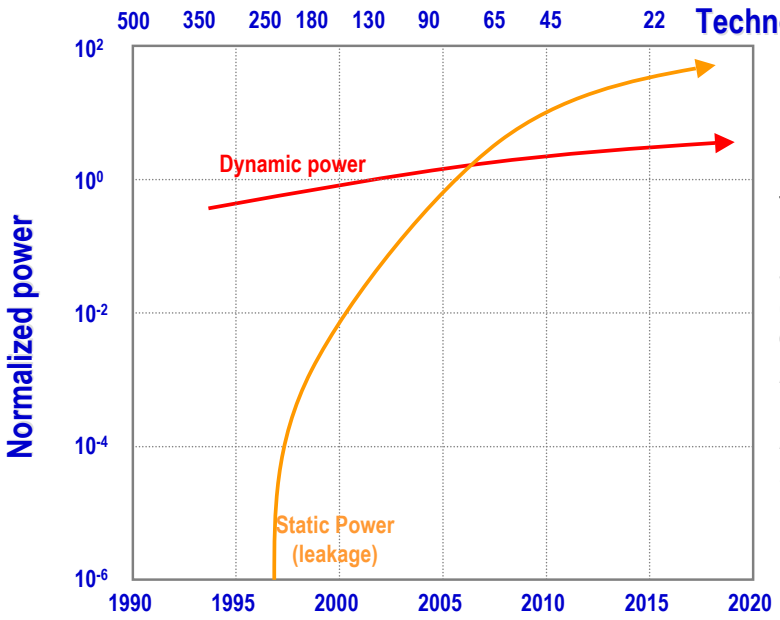

International Technology Roadmap for Semiconductors (ITRS)2001, 2002. Courtesy: *Moore's Law Meets Static Power, Computer, December 2003, IEEE Computer Society*

Static Power can surpass Dynamic Power

Moore's Law…

The Good…

More logic, smaller, faster, cheaper transistors = Progress

the Bad…

More logic + thinner oxide, more leakage + higher clock rate =
More power, shorter battery life, more heat to be dissipated

and the Ugly:

More heat = higher temperature, lower performance, shorter life

Arrhenius model:

Half the life expectancy for every 10°C temperature increase

The Near Future (2010)

- Technology will evolve to 45nm, then 32nm
  - Half the cost per function, compared to Virtex-5
  - Twice the high-end capacity, compared to Virtex-5
- Vcc probably lowered (to 0.7 V ?)
  - Reduces dynamic and static power consumption
  - Delicate balance between speed & leakage current
- Intensive search for low-K dielectric
  - To reduce interconnect capacitance
- Intensive search for high-K gate-dielectric (hafnium?)
  - high transconductance + low leakage
- Power consumption is biggest concern
  - must be kept within reasonable limits

Moore's Law is still alive,
but higher clock speed gets harder to achieve

Conclusion

- FPGAs are gaining popularity,
  - No more "too small, too slow and too expensive"
  - Ideal for fast-changing standards and rapid product cycles
  - ASICs are only for very high volume, speed, or ultra-low power
- Hard cores improve performance and ease-of-use
  - massive parallelism (up to 200,000 six-input LUTs)
  - outperforms even the fastest dedicated DSP chips
- FPGAs benefit from the cutting edge of Moore's law,
  - but higher performance comes mainly from better architecture