Technology and Market Trends

„a pseudo-random walk“

The Technology defines what is feasible in computing

The Market defines what is affordable in computing

The infrastructure boundary conditions define what is implementable in computing

Some background

These CERN technology and market investigations are done in a general way, but need to take into account the boundary conditions coming from the High Energy Physics community and the CERN infrastructure.

→ Cost predictions for budget and resource planning

Basic features:

-- program performance is determined by integer calculations (80%)
-- events are independent, thus no fine grain parallelism is needed
-- programs need >= 2GB of memory per job (=core)

Computer Center:

-- 10000 low-end servers installed
(dual CPU, >= 2GB memory per core, 1-6 TB local disks, 80% 1Gbit – 20% 10 Gbit, 24-36 disk in internal or external data disk trays)
-- 65000 cores, 62 PByte raw disk capacity, 50 PB data on tape
-- 3.5 MW power and cooling

-- replacement and purchase rate is about 1500 servers per year (in AND out)

Semiconductor Industry

Worldwide Revenue Ranking for the Top-25 Semiconductor Suppliers in 2011
(Revenue in Millions of U.S. Dollars)

2010 Rank	2011 Rank	Company Name	2010 Revenue	2011 Revenue	Percent Change	Percent of Total	Cumulative Percent
1	1	Intel	40,394	48,721	20.6%	15.6%	15.6%
2	2	Samsung Electronics	28,380	28,563	0.6%	9.2%	24.8%
4	3	Texas Instruments	12,994	13,967	7.5%	4.5%	29.3%
3	4	Toshiba	13,010	12,729	-2.2%	4.1%	33.4%
5	5	Renesas Electronics Corporation	11,893	10,648	-10.5%	3.4%	36.8%
9	6	Qualcomm	7,204	10,198	41.6%	3.3%	40.1%
7	7	STMicroelectronics	10,346	9,735	-5.9%	3.1%	43.2%
6	8	Hynix	10,380	9,293	-10.5%	3.0%	46.2%
8	9	Micron Technology	8,876	7,365	-17.0%	2.4%	48.6%
10	10	Broadcom	6,682	7,160	7.2%	2.3%	50.9%
12	11	Advanced Micro Devices (AMD)	6,345	6,436	1.4%	2.1%	52.9%
13	12	Infineon Technologies	6,319	5,312	-15.9%	1.7%	54.6%
14	13	Sony	5,224	5,015	-4.0%	1.6%	56.3%
16	14	Freescale Semiconductor	4,357	4,408	1.2%	1.4%	57.7%
11	15	Elpida Memory	6,446	3,887	-39.7%	1.2%	58.9%
17	16	NXP	4,028	3,831	-4.9%	1.2%	60.1%
20	17	nVidia	3,196	3,608	12.9%	1.2%	61.3%
26	18	ON Semiconductor	2,291	3,428	49.6%	1.1%	62.4%
18	19	Marvell Technology Group	3,606	3,393	-5.9%	1.1%	63.5%
15	20	Panasonic Corporation	4,946	3,390	-31.5%	1.1%	64.6%
21	21	ROHM Semiconductor	3,118	3,187	2.2%	1.0%	65.6%
19	22	MediaTek	3,553	2,952	-16.9%	0.9%	66.6%
28	23	Nichia	2,190	2,936	34.1%	0.9%	67.5%
22	24	Analog Devices	2,862	2,846	-0.6%	0.9%	68.4%
23	25	Fujitsu Semiconductor Limited	2,757	2,742	-0.5%	0.9%	69.3%
		All Others	96,073	95,610	-0.5%	30.7%	
		Total Semiconductor	307,470	311,360	1.3%	100.0%	

Source: IHS iSuppli March 2012

50% of the market is shared by only 10 companies

INTEL is the largest company

This covers:
Processors
Memory
Flash memory
Microcontroller
Components
……

Worldwide semiconductor market revenues: 311 B$ in 2011
3-4 % growth rate expected for 2012

Computing Market in 2011

Mobile Devices

Smartphone: 490M, +63%, 160 B$

Tablet: 63M, +274%, 26B$

Personal Computer

PC: 351M, +3%, 260 B$

- Desktop: 112M, +2.3%

- Notebook: 210M, +7.5%

Ultrabook: 4M

- Netbook: 29M, -25%

Server

Server: 9M, +8%, 52 B$

- HPC: 10 B$, +8%

Supercomputer:
4.4 B$, +8.5%

Type: number of units shipped, revenue growth rate, market revenue

5-10% variations: IDC, Gartner, iSupply, etc.

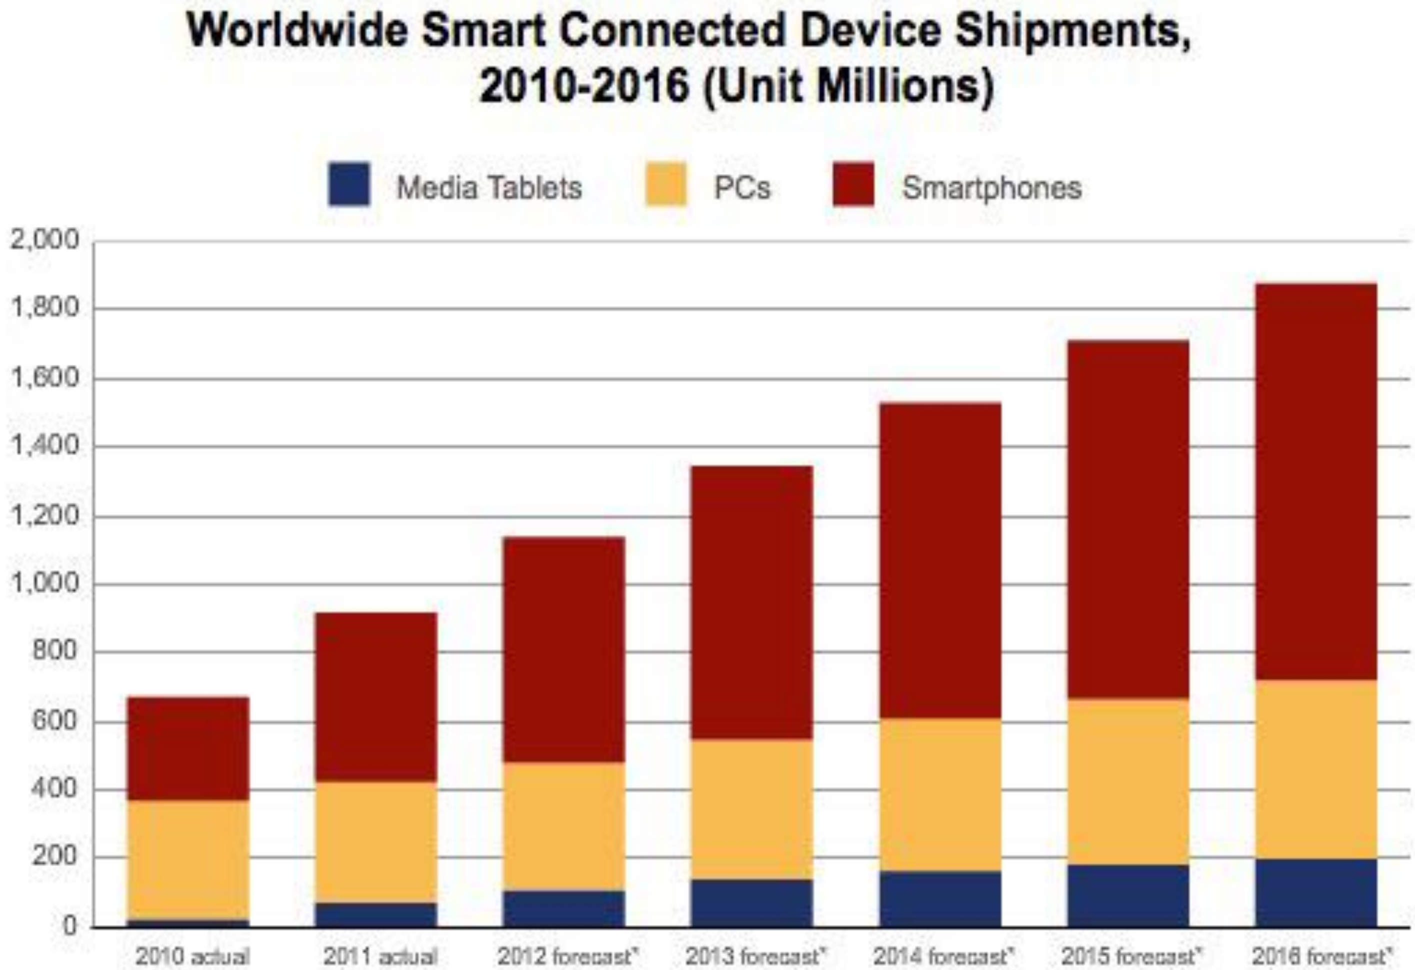

IDC predictions

In April 2002 the 1 billionth PC was shipped (PC = desktop+notebook+server)
2 billion until 2007 and about 3.5 billion right now
Compared to an estimated 1.5 billion installed PCs today (2 billion in 2015)

~6 billion mobile phone subscriptions worldwide at the end of 2011

Computing Market Trends

- The number of sold Smartphones is now higher than the amount of PCs sold
- Smartphones and tablets have the highest growth rates
  → Pushes the combination of mobile devices and cloud computing
- The netbook will disappear
- Ultabooks are a new category; INTEL estimate for 2012 : 20-30 M units?!
- The server market has still healthy revenue growth rates,
  but last year the prices increased and less units were shipped
  → High Profits and Cost (HPC and Supercomputer)
  → Consolidation and efficiency improvements (virtualization) in large computer centres is affecting the server market

Market push from raw performance to power efficiency (execution and standby)
→ few cores, simple processors, specialized processors (DSP, FPGA), SSD disks, NAND flash memory,…… **"Dark Silicon"**

Processor Technology

Main focus is to reduce leakage currents
and reduce voltages
→ Energy efficiency

Add more functional units on the die
(cores, GPUs, memory, IO, etc.)
but keep surface area constant

3D tri-gate transistors, FinFET

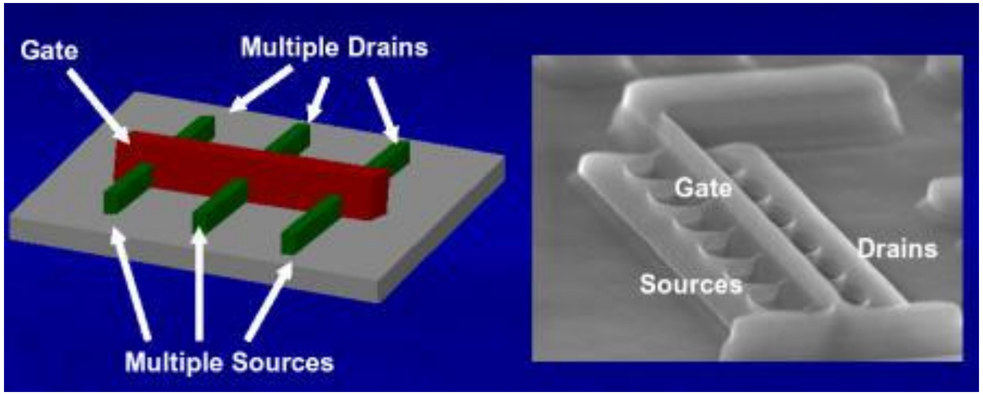

Special technology necessarry for each shrinking step:

- •90nm strained silicon
- •45nm high-k metal gates
- •22nm 3D tri-gate transistors

High-k metal gate transistors are strained, and FinFET transistors have both strained silicon and high-k metal gates

INTEL has produced the first 14nm samples in the lab

But the current optical immersion-lithography (193nm) is in principle not capable of producing 14nm structures
extreme ultraviolet lithography is needed → very complicated, expensive, not all problems solved yet, 2015 target will be tight

Intel Server Microarchitecture Roadmap

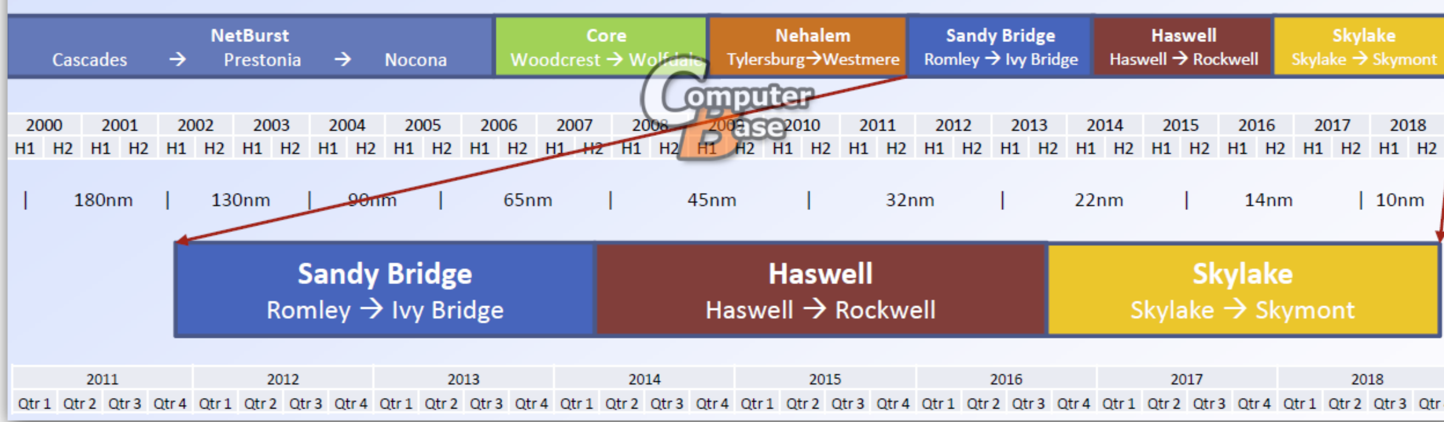

Intel 12-18 month ahead of the competition: 22nm Ivy Bridge processors released
AMD and Nvidia had/have some problem with their 28nm processes
(yield and general processing with TSMC)

A Fab for 10nm structure production will cost ~10 B$; only very few companies will be able to effort this

Applied Material, company in the ‘background’ provides tools, machinery and expertise for chipmakers;   INTEL and AMD (TSMC, GlobalFoundries)

Multi-Core I

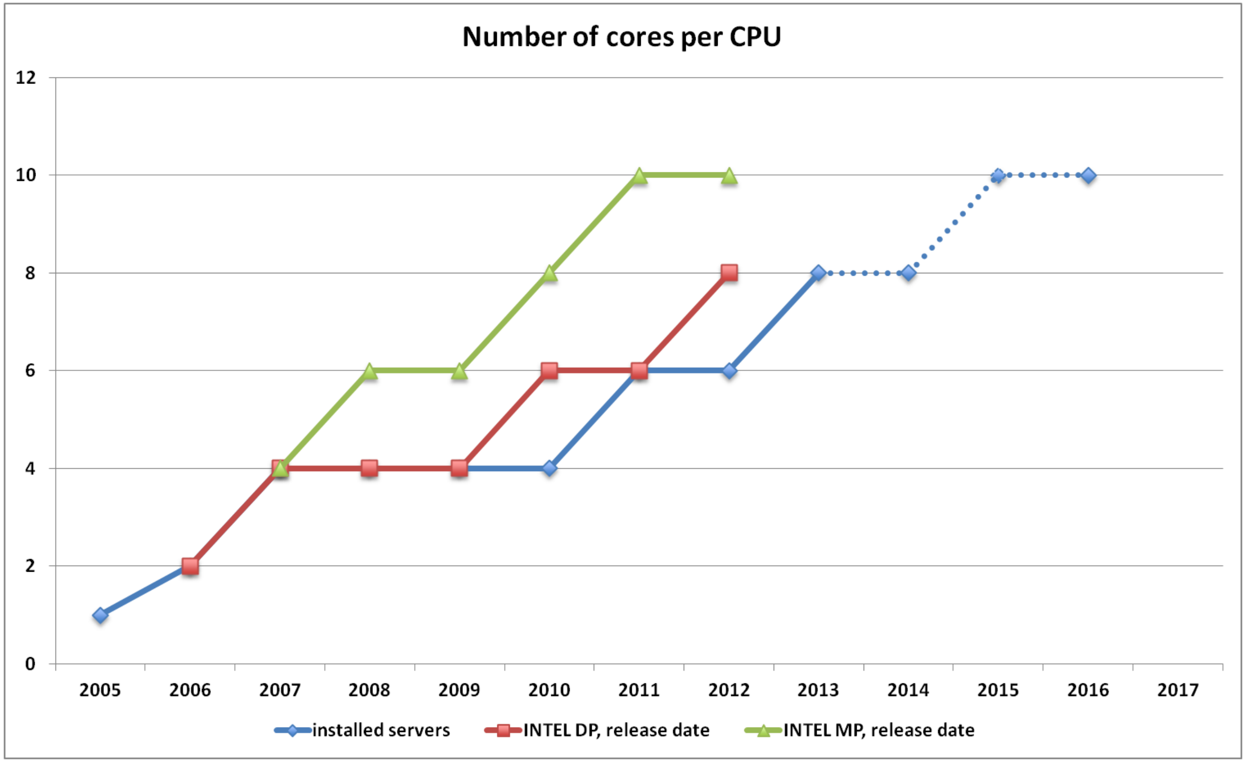

Linear increase, about +2 every second year

Processors installed at CERN and general INTEL server processors (max cores at release date)

High end processors from the top-500 end 2011: Sun Niagara2: 8-cores, IBM Power7: 8-cores, Fujitsu Venus: 16-cores, AMD Interlagos: 16-cores

Important factor: memory per core, number of IO streams ("disk-spindles per core")

Multi-Core II

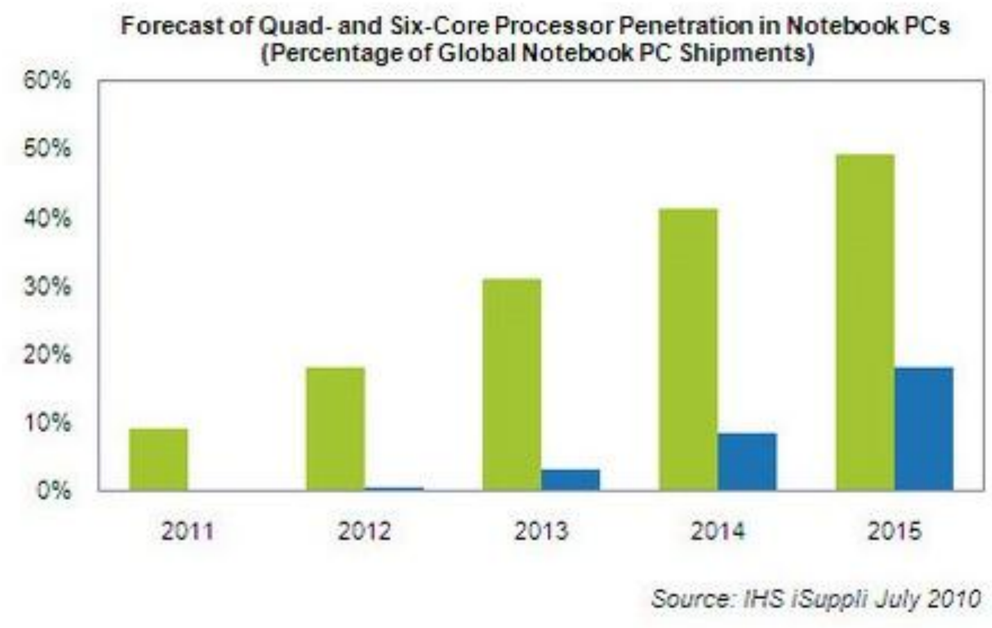

Slow increase in core count for the Mobile, PC and low end server market

First quad core processors in the Smartphone market, mid/end 2012
Low volumes

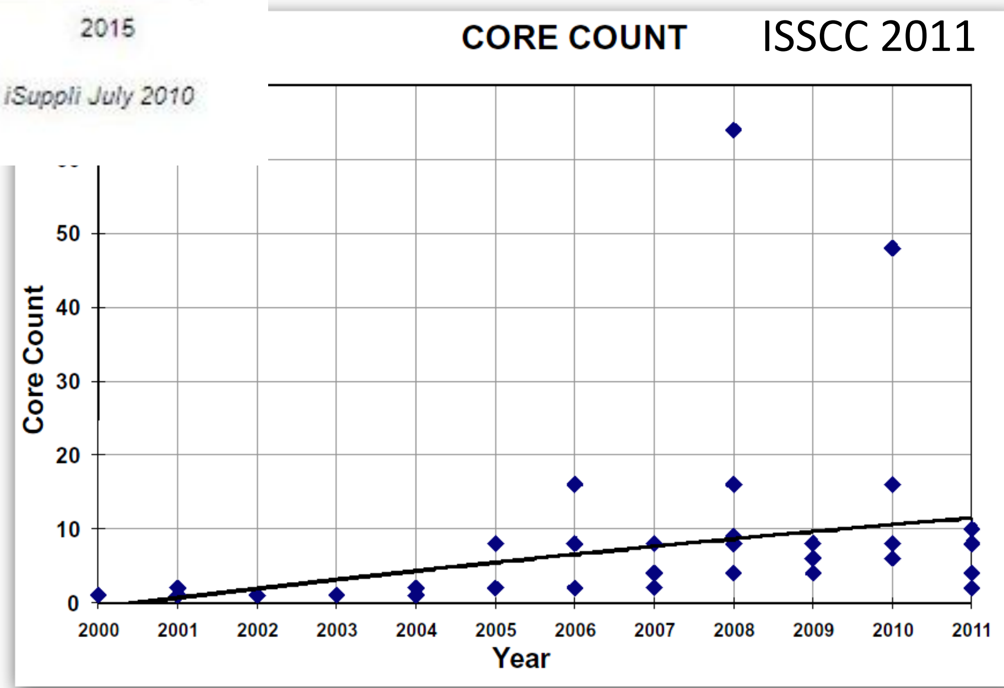

ISSCC
International Solid-State
Circuits Conference Trend Report 2011

Many Cores

Tilera: 64bit processors in a two dimensional array; multiple mesh networks; Just released 36-core 64bit processor version;
Optimized for key-value access → Facebook tests

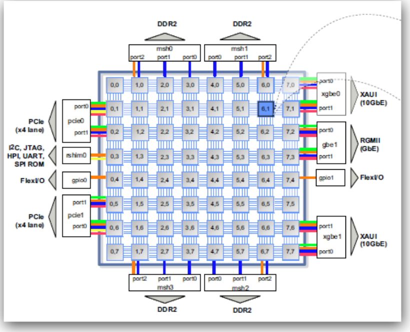

INTEL: MIC Knights Corner, 50-60 cores; 48-core SCC 'cloud computer'

Calxeda: servers based on ARM Cortex-A9
Used by HP (Project Moonshot) , 288 processors in 4U

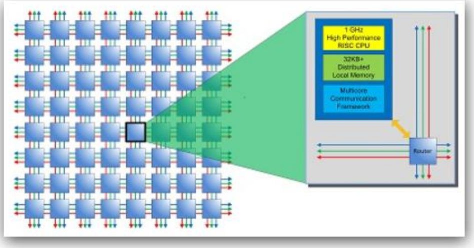

Adapteva: Epiphany IV processor, 64 RISC processors

Enterpoint: PCIe co-processor board, Xilinx FPGA

ZiiLabs: 100-core system, Cortex-A9 + 96 StemCell Media Processing cores

Plurality: Hypercore processor, 256-cores, e.g wireless ingfrastructure

Neuromorphic processors e.g. silicon retina, ‚Third Eye' motion detection, SyNAPSE project using memristors

What about: Larrabee, Cell processor, Kalray, ...

Chip Complexity

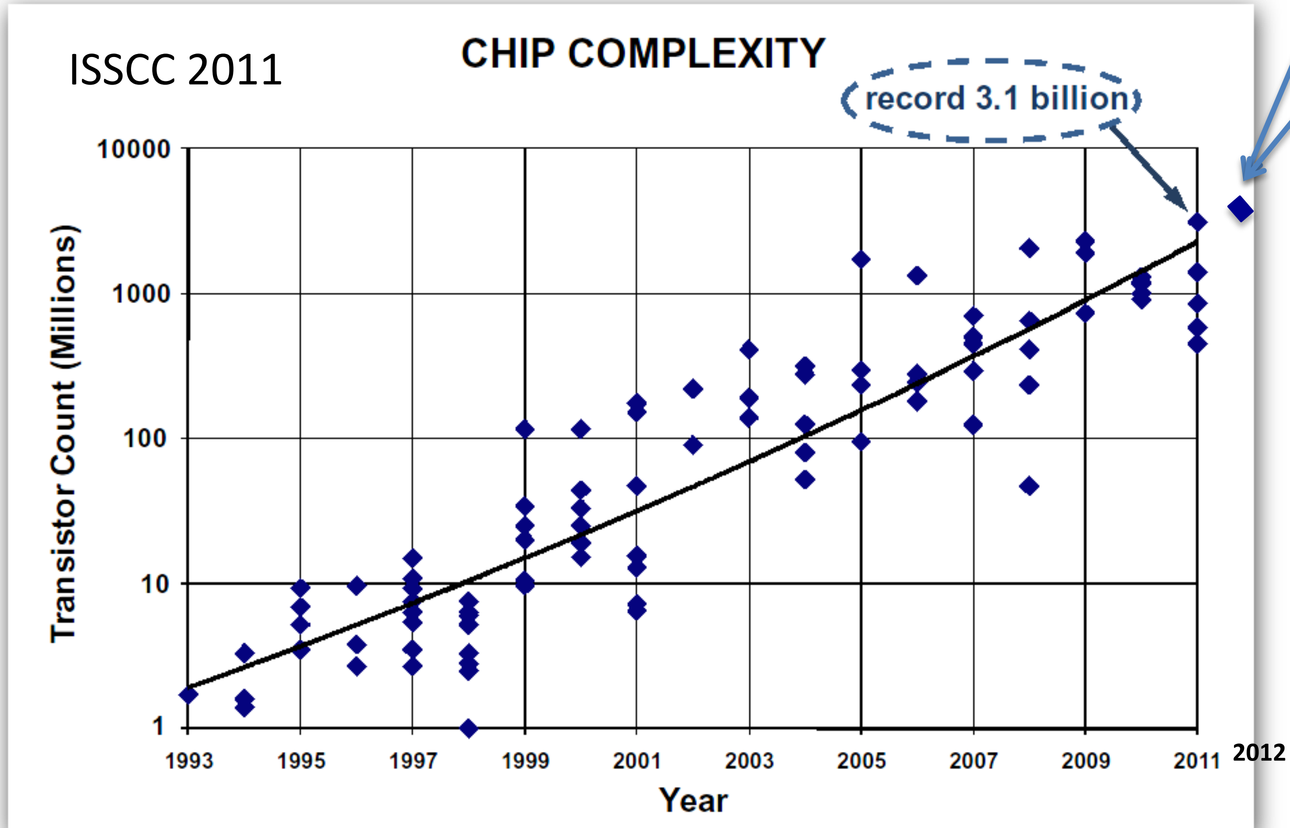

GPUs

4.3 billion, Radeon 7970

3.5 billion, Nvidia 680 GTX

Still following Moore's Law

One has to take into account that on the processor chips about 50% of the Transistors are used for L1+L2+L3 caches and the rest is distributed across the cores.

INTEL/AMD core : 100 million transistors ARM core: 6 million transistors

Future Transistor Technology

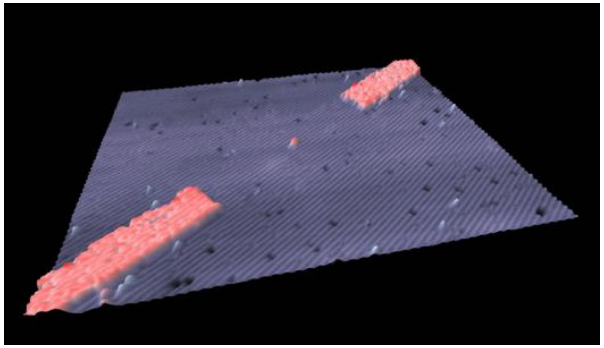

Single phosphorus atom transistor on silicon

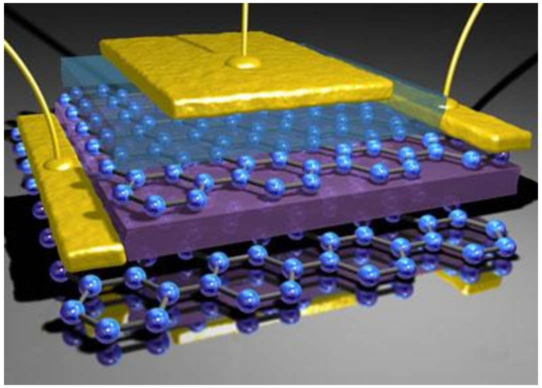

Single atom layered transistors based on graphene and boron nitride

9nm transistor based on nanotubes

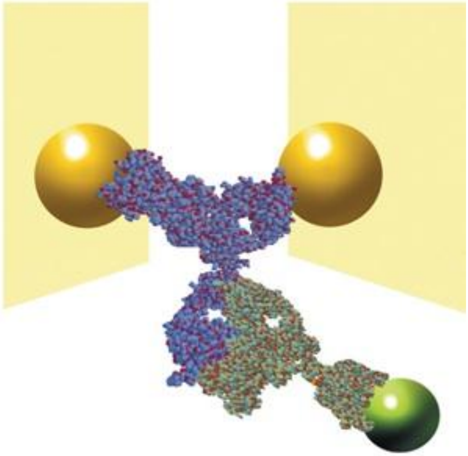

Self-assembly of transistors based on biological molecules (immunoglobulin G antibody) and 5nm gold particles

Very active research area; frequently new results published, but will not effect the market in the <5 years time frame

From components to server costs

Cost calculations for a dual-processor server node:

- Processor cost is the dominant factor
- Fluctuating memory costs
- Server motherboard + chassis + power supply, relative constant costs
- Local disk space (HDD or SSD)

Boundary conditions

- Changing requirements from applications
  more memory, move from 1 to 10 Gbit network, more
  and faster local disk space
- ‘rounding’ - matching memory-per-core with number of memory DIMM slots on the motherboard, memory channels, DIMM size, power requirements, memory costs
- Power efficiency of the overall system
- Computer center limits:
  ceiling on power and cooling, power and cooling density, space

Processors Costs

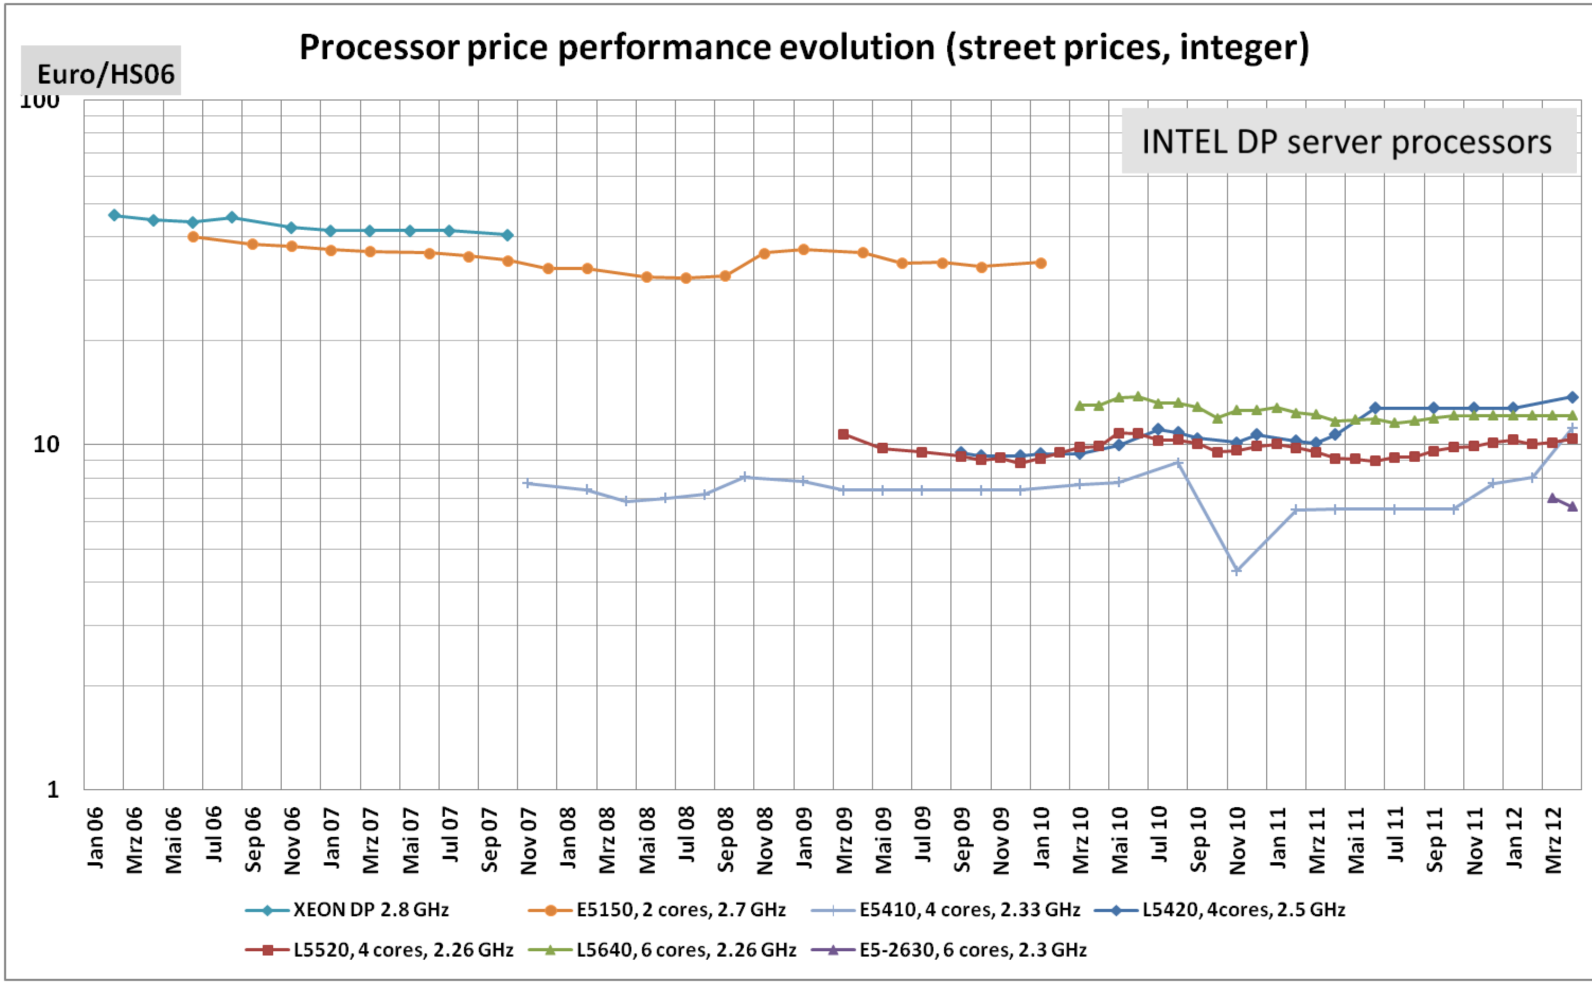

HS06= HEPSpec2006, integer SPEC benchmark

NO price improvements per processor
Step function for different generations

Ratio specint/specfloat:
3.5 +-0.5 in 2006-2010
2.4 in 2011

Memory Costs

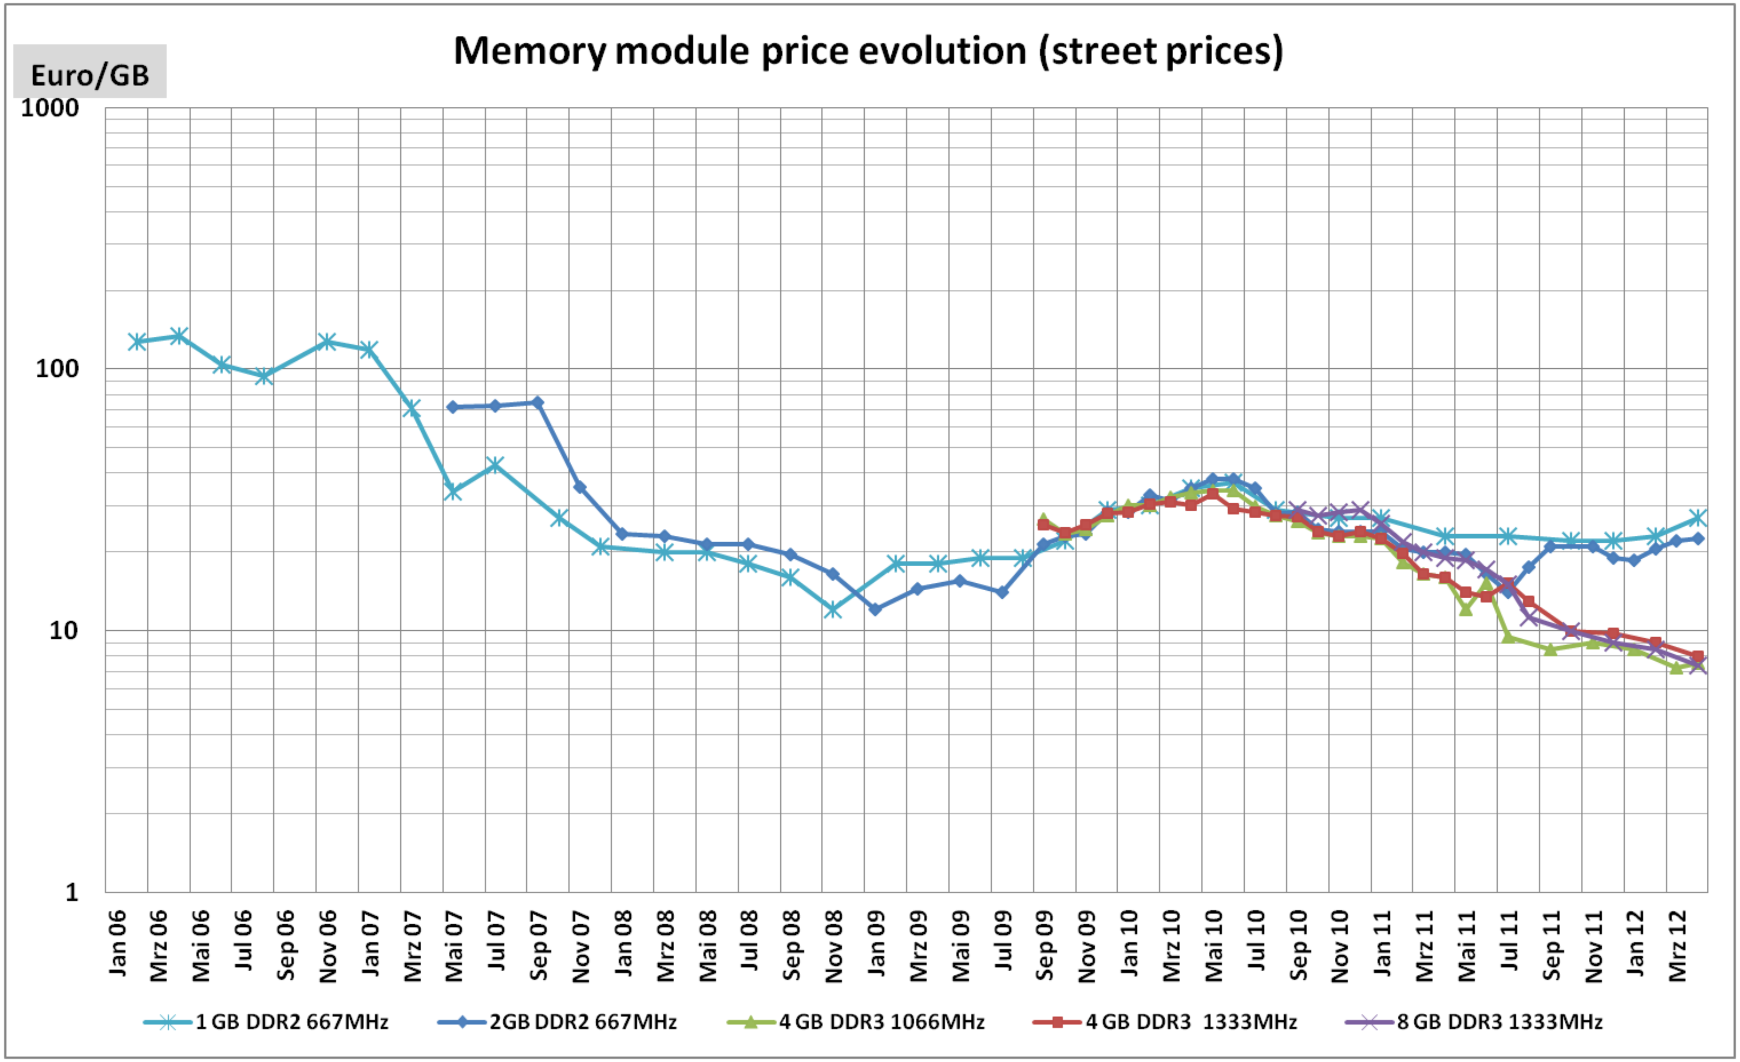

ECC registered server quality memory

Large fluctuations, volatile market coupled to NAND flash memory production

Server Costs

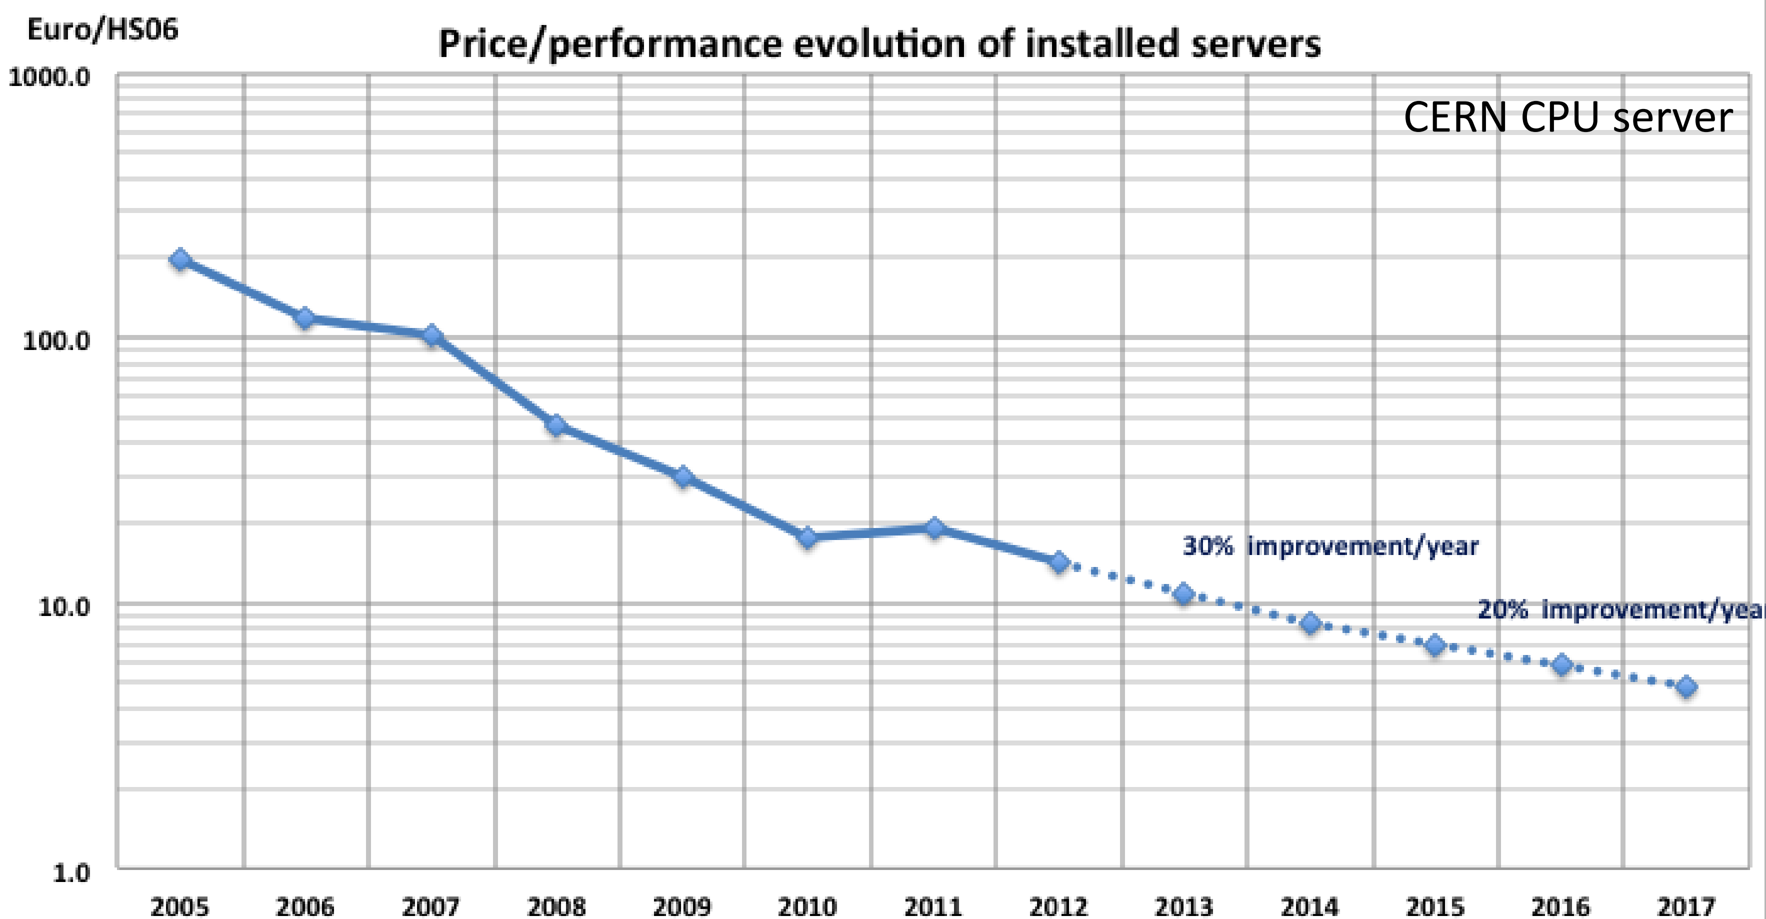

Ratio of server costs over processor costs → 2006-2009 : 2.7 2010-2012: 1.7
= Processor cost share of a server went from ~40% to 60%

Expect a decrease in improvement speed when the technology reaches 14 nm structure sizes

Storage

DRAM Memory

- 4 companies have 91% market share: Samsung, Hynix, Micron, Elpida (bancrupt)
- 800 million units sold == 2 ExaBytes

NAND Memory

- 4 companies have 99% market share: Samsung, Toshiba, Micron, Hynix
- 4000 million units sold == 19 ExaBytes

Solid-State-Disks

- > 50 companies
- 17 million units sold == 3 ExaBytes (included in the NAND memory numbers)

Hard-Disk-Drive

- 3 companies only: Western Digital 50%, Seagate 39%,Toshiba 11 %
- 630 million units sold == 330 ExaBytes

Magnetic Tape

- 3 technologies : IBM, Oracle, LTO-consortium LTO has 90% market share
- 27 million units sold == 20 Exabytes

Reference year is 2011

Storage Feature Size Predictions

Ref.: IBM April 2012 RFontana

Lithography Roadmaps

- **Minimum feature typically reduced by 12% per year**
- **Intel/Micron has consistently exceeded ITRS goals**

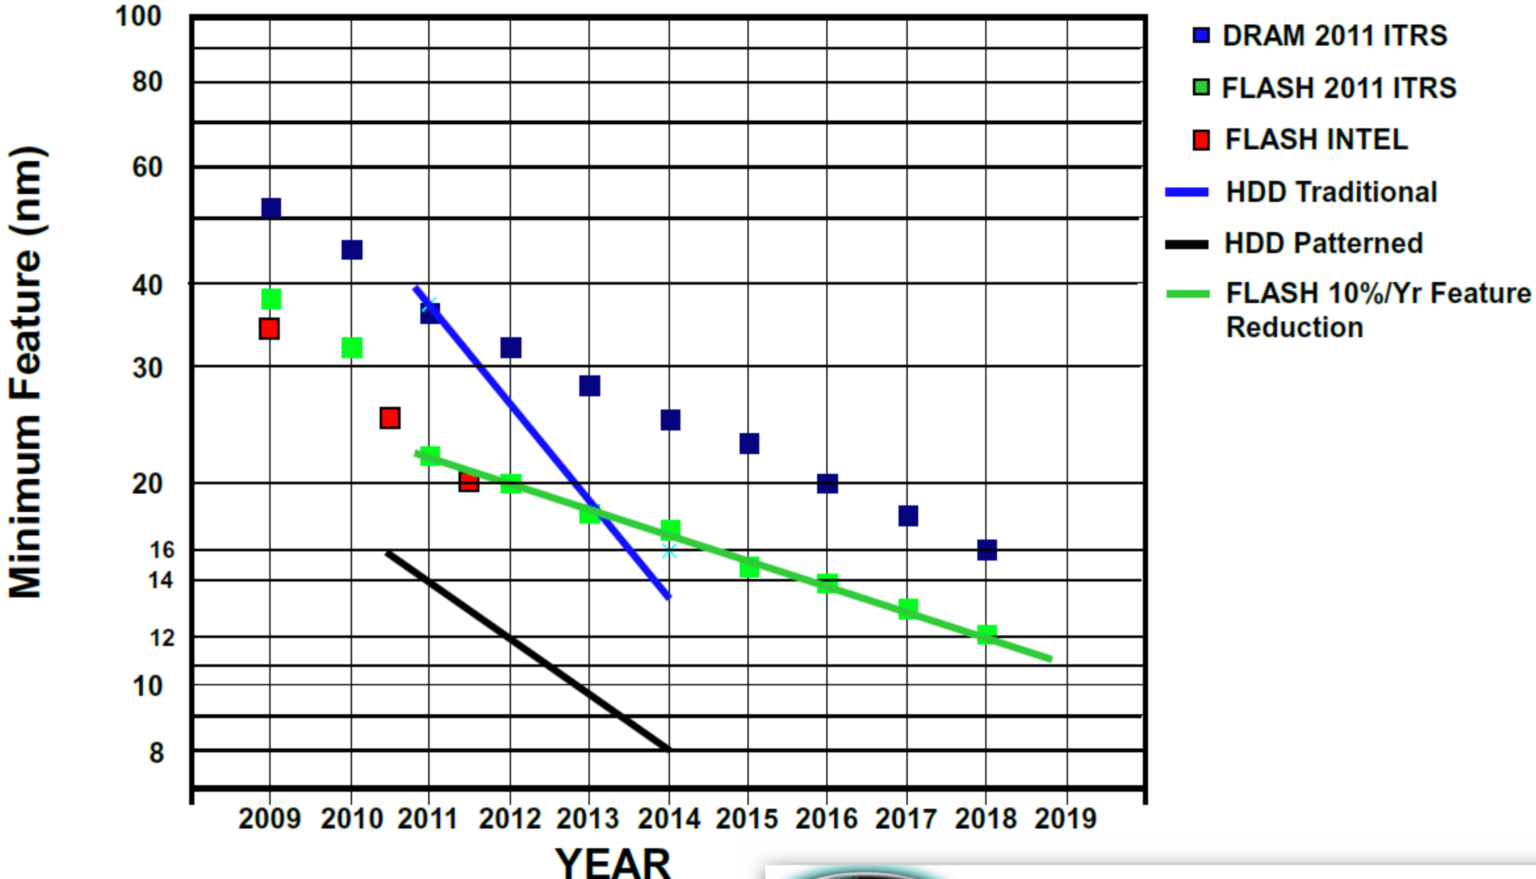

International Technology Roadmap for Semiconductors

Storage Areal Density Evolution

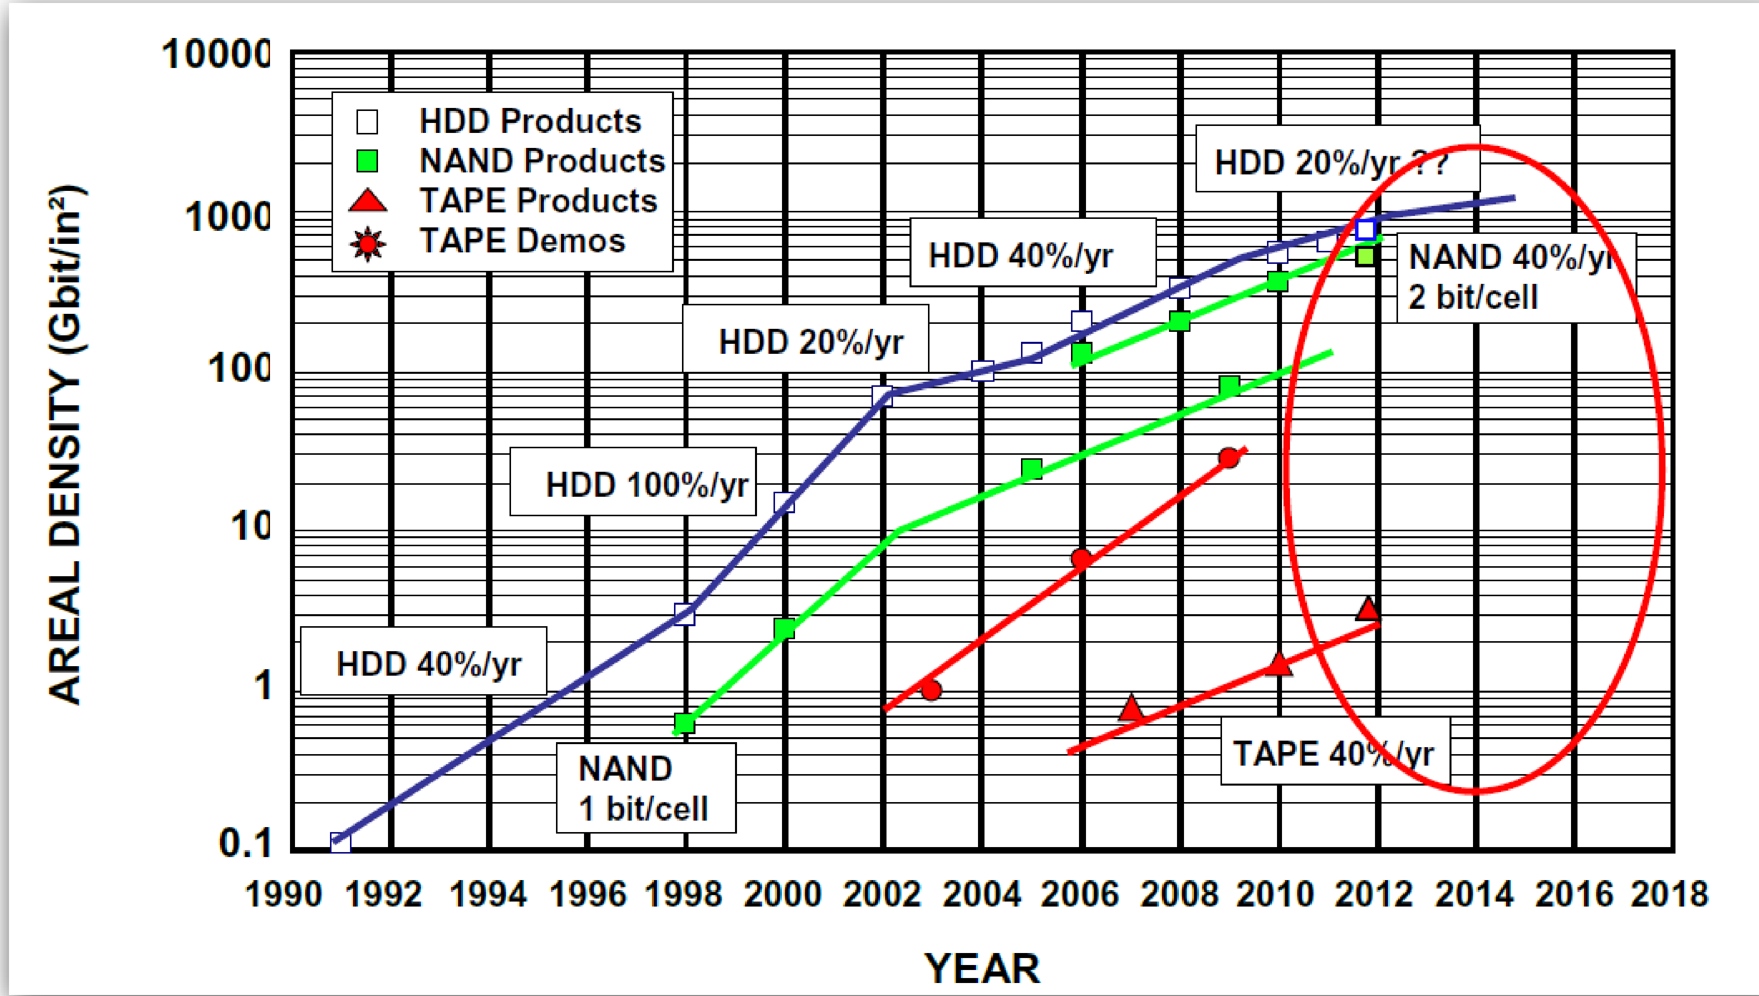

Ref.: IBM April 2012 RFontana

Storage Areal Density Predictions

Annual Areal Density Growth Rate Scenarios

- HDD – 20% to 25% – Transition to New Technology, Sensor Output, Lithog
- NAND Flash – 25% to 30% – Lithography and Endurance
- TAPE – 40% to 80% -- No Lithography Issues, Mechanical Realities

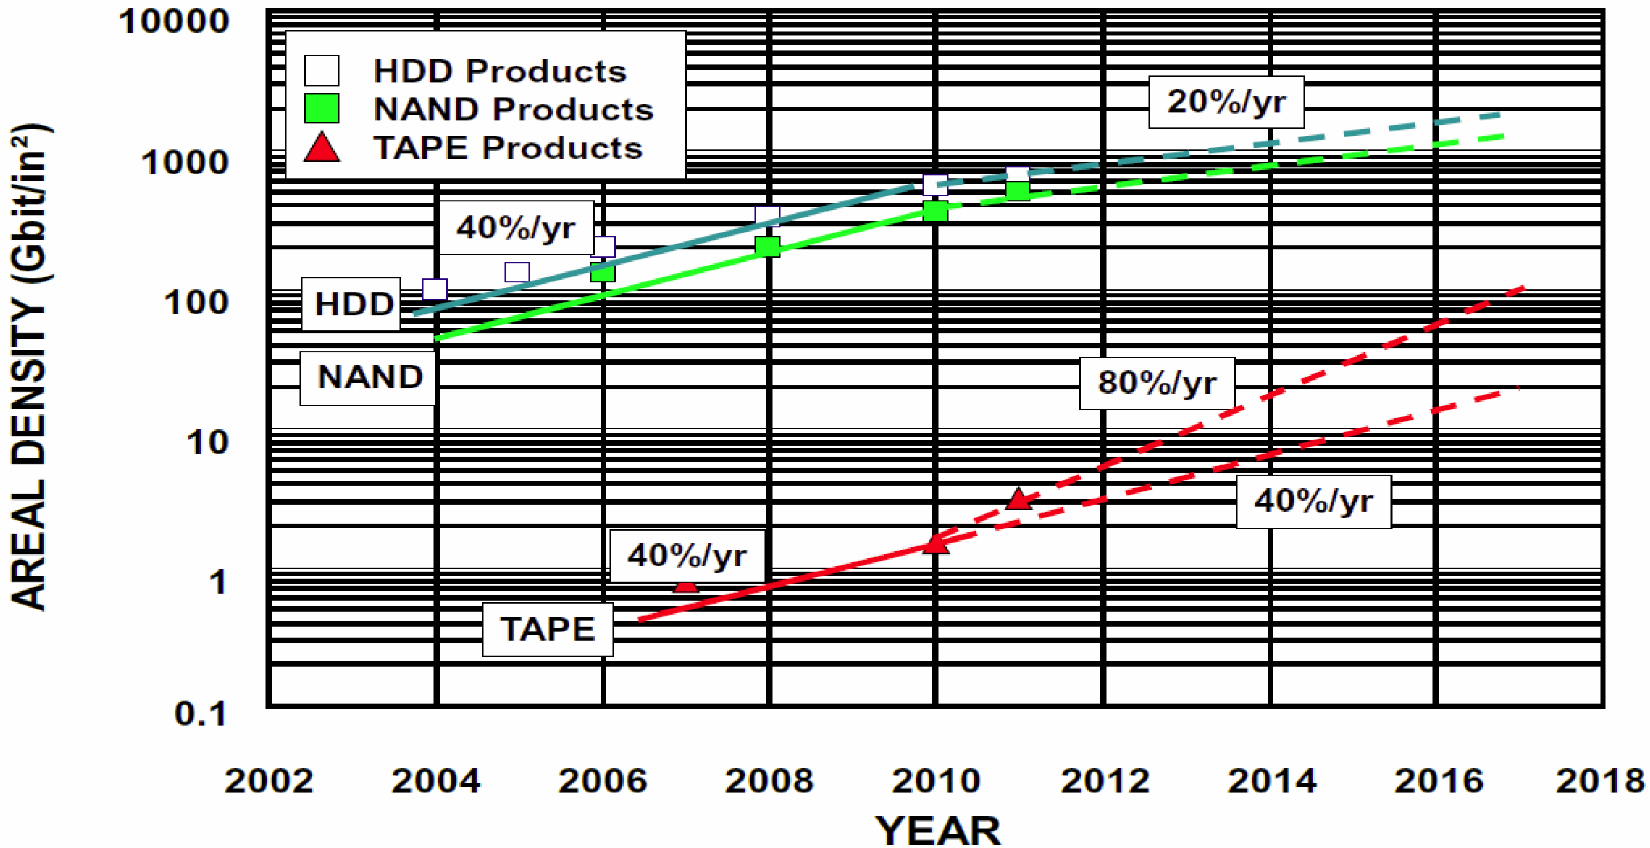

Ref.: IBM April 2012 RFontana

DRAM Memory

- DRAM market revenue was 31 B$ in 2011
- Revenues dropped by -24% in 2011→ PC market changes, focus on NAND
- Moving to DDR4 (1866 MHz) in 2013, but focused on the high end server market
- Roadmaps to move to 3D transistors and hyper memory cube
- Problems to improve bandwidth and power consumption
  → multi-core and multi-threading memory bandwidth demand
- Scaling problems in the 2x nm range (Lithography)

More price fluctuations expected

Jedec DRAM Memory Roadmap

	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020
Process	3x nm		2x_H		2x_L		1x_H		1x_M	
DDR3	1600		1866							
DDR3L	1333		1600							
DDR4			1866	2133	2400	2667				
DDR4L						2400	2667	2667	2932	3200
Device		2Gb								
Device				4Gb						
Device							8Gb			
Device										16Gb
DIMM	8GB	16GB		32GB		64GB	64GB			128GB
3DS/TSV				DDR4_2H	DDR4_4H		DDR4_8H			

Note:
* DRAM speed: device raw speed, in Mbps.
* DIMM density: sweet spot density.
* 3x = 30-39 nm, 2xH = high 20's nm, 1xH = high teen nm, 1xM= mid teen nm
* DDR4L: 1.0V, TBD.

Jedec :Joint Electron Device Engineering Council
→developing open standards for the microelectronics industry

Hard Disk Drive

- All drives today use perpendicular magnetic recording technology (PMR)
  Introduced in 2006

- Highest density in production: 740 Gbit/in^2 for 2.5” disk and 620 Gbit/in^2 for 3.5” disks
  1 TBbyte per platter

- PMR reaches its limit at about 1 Tbit/in2

- Future technology is heat-assisted magnetic recording (HAMR)
  Seagate demonstrated 1 Tbit/in^2 density

- HAMR is limit to 5-10 Tbit/in2 → 30-60 Tbyte for a 3.5” disk
  large production level probably only in 2016-2017
  very expensive technology
  → areal density growth rate slows down to 20-25%

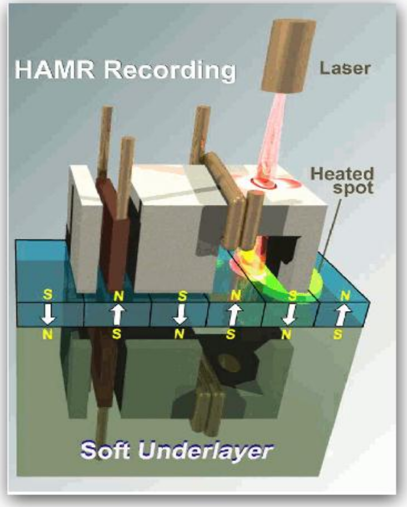

HDD space demand comes 50% from cloud storage
2011 about 400 ExaBytes needed , but only 300 Exabytes delivered
The demand will increase to one Zettabyte in 2016 (estimate)
→ Investment into production capacity: larger AND more disks to be shipped
Higher demand than production rate = price decrease slowdown

Hard Disk Drive Costs I

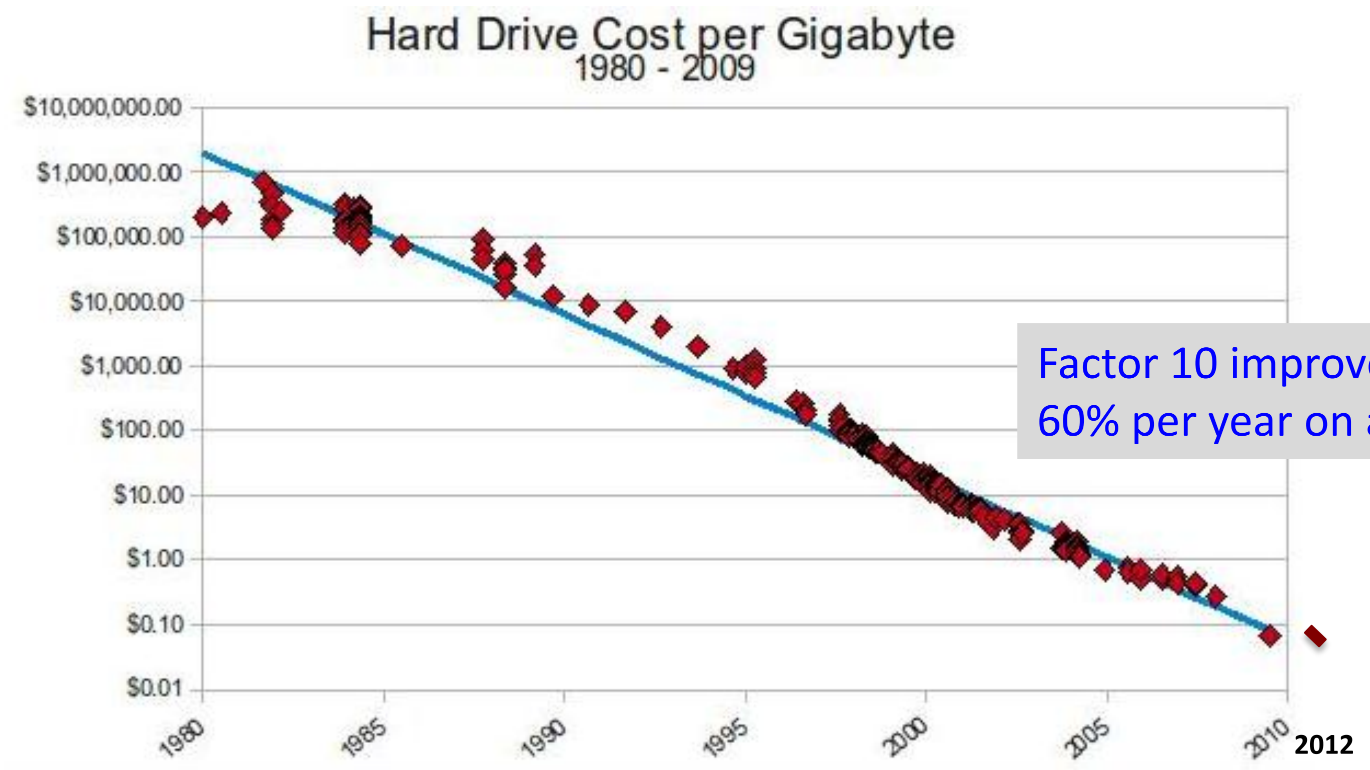

Factor 10 improvement in 5 years = 60% per year on average

Desktop disks: up to a factor 2 cheaper than server disks, not qualified for 24*7 uptime
→ Low MTBF !? Would be interesting to have a large scale investigation…….

Hard Disk Drive Costs II

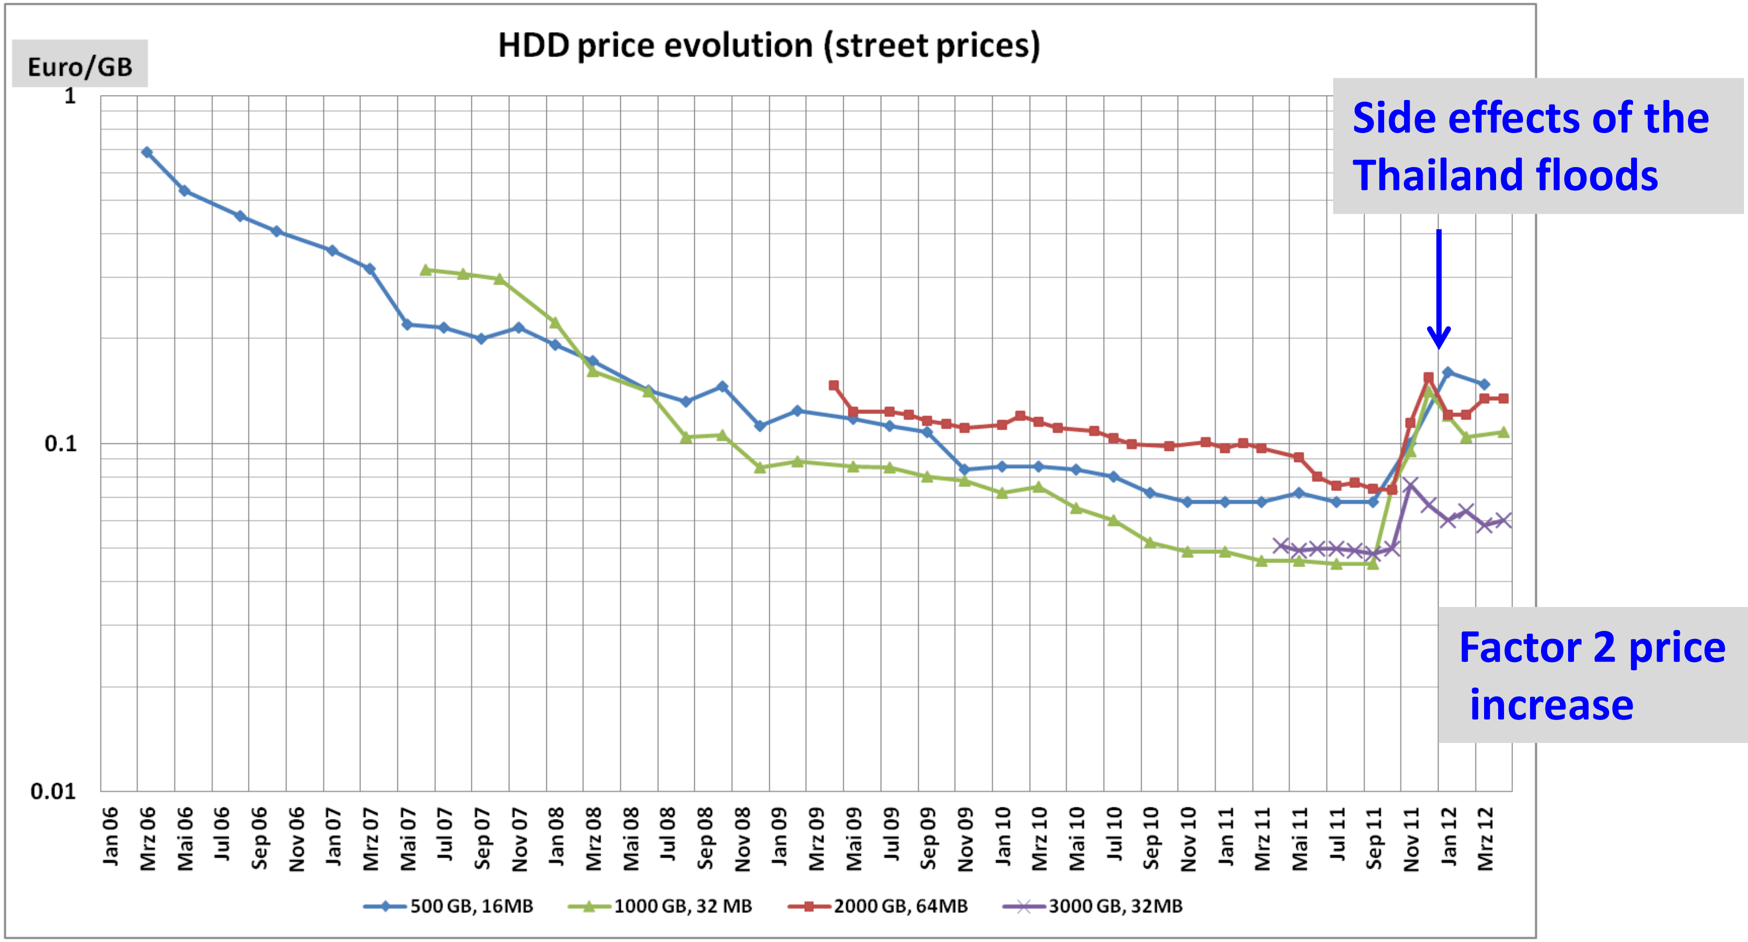

Server quality HDDs: 3 years warranty, SATAT II, large cache, 24*7 operation, high MTBF

(we actually measure 300000h while the vendor claims 1-1.5 million hours MTBF)

NAND and Solid State Disks

NAND flash market ~30 B$, expectation is 10% growth rates
(smartphones, tablets, ultrabooks)

- Today: 64Gbit in 25-28nm structure sizes, 2bit cells
- SanDisk+Toshiba prototype: 128Gbit, 19nm, 3bit MLC

Lithography problems below 20nm
Reduced endurance levels

More bits per cell + smaller structure sizes
= better GB/$ and worse endurance

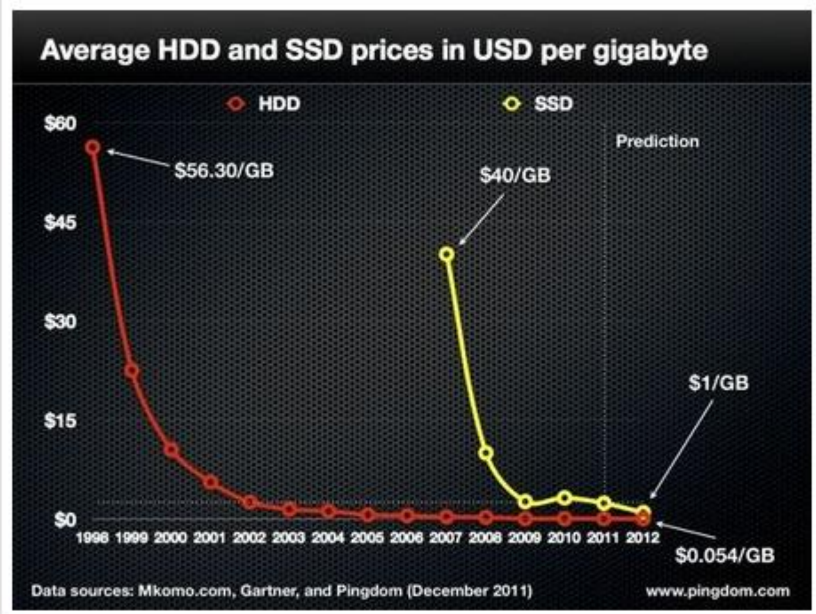

Solid-State-Disks

- SSD growth rate very high: from 17m units 2011 to an expected 45m in 2012
- Lowest cost is 0.8 Eur/GB, but up to factor 20 variation in cost
  factor 10-50 compared to HDD costs

For SSDs Euro/GB is actually the ‘wrong’ metric → IOPS and sequential performance
SSD cache + HDD backend

Future Storage Technologies

Trying to combine the advantages of HDD, Flash and DRAM
Non-volatile, Fast read and write, high endurance, High 2D or 3D density

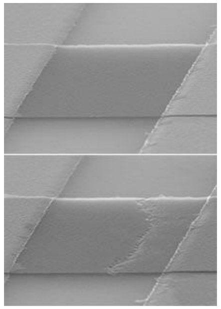

IBM: racetrack memory
moving domainwalls in a nanowire

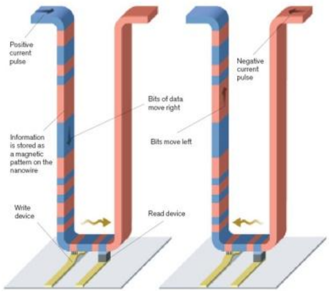

Rice-University: 3D graphene storage
based on creating cracks

- Samsung: combing graphene and silicon
- University Singapure: combining graphene with ferroelectrics
- University California/Taiwan: embedded 3nm silicon nanodots

Memristor first implementation expected in mid-2013
2008 predictions: replacing Flash in 2012, DRAM in 2014, HDD in 2016 !

PCRAM Phase change memory: 8Gbit 20nm PCR prototype device by Samsung in Nov 2011

MRAM already in production state, 16Mbit chips, 130nm structure size, non-volatile, 50ns latency, ~3 million units shipped in 2011

Many innovative technologies, but no market relevance in the 3-5 years time frame

ITRS technology assessment Nov 2011

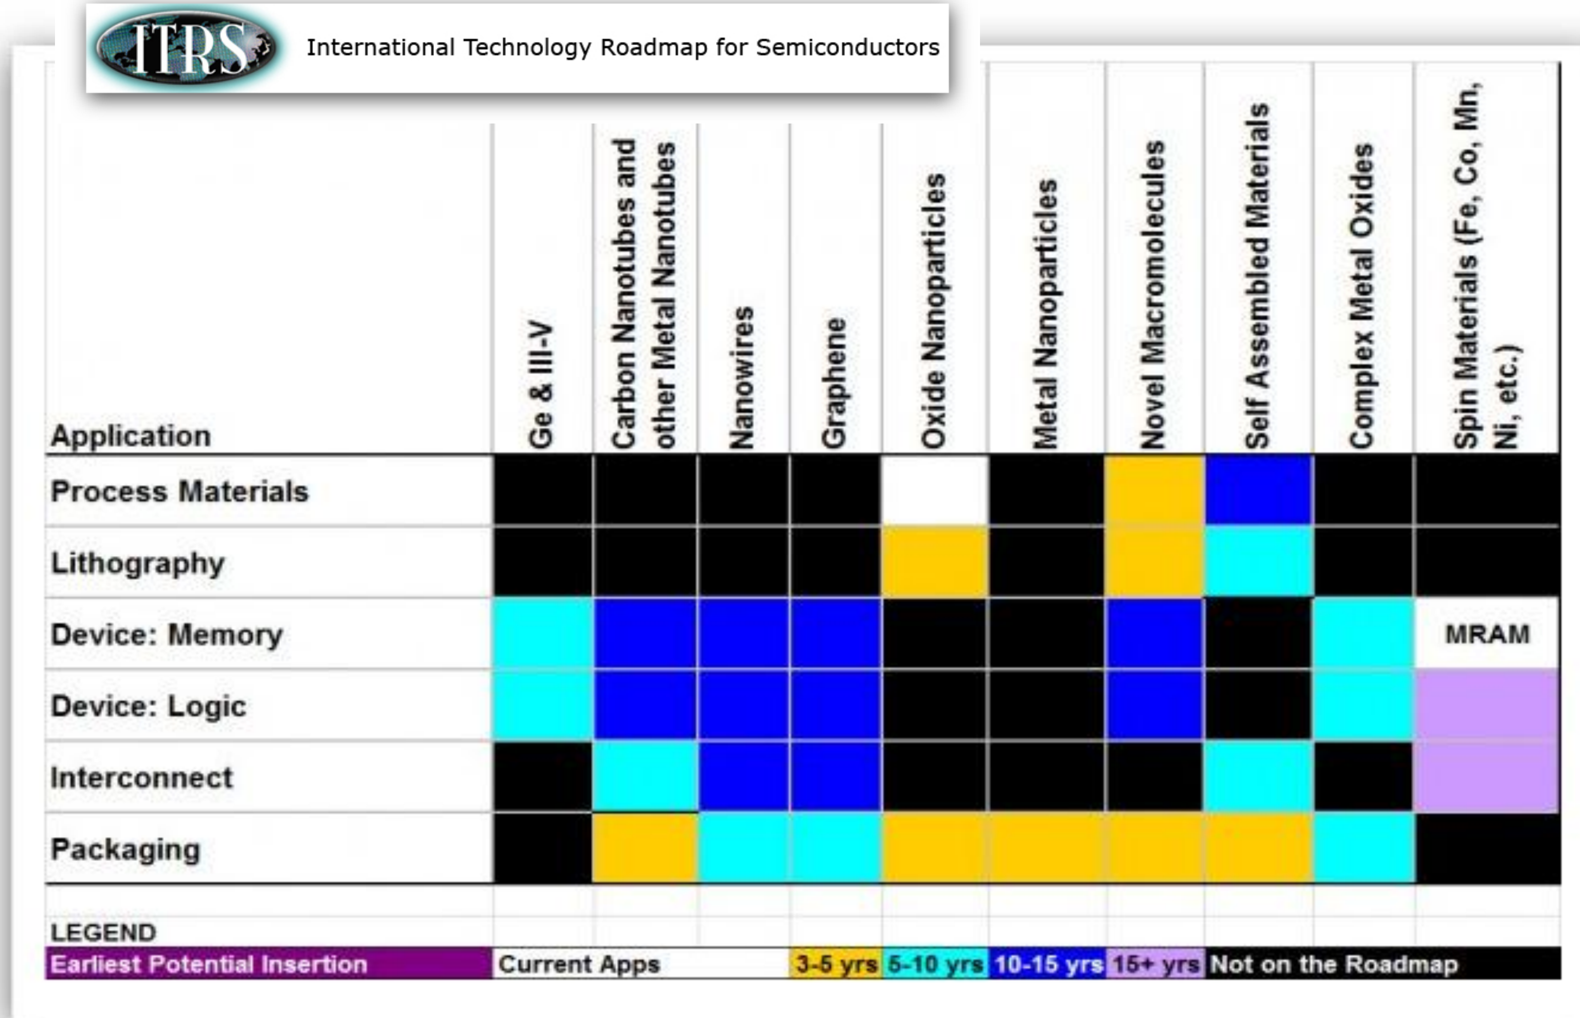

Adiabatic versus disruptive technology changes

What happened to: holographic storage, Tesa-ROM, millipede, ………

Graphics

Separation of discrete graphics card, on-board GPUs and CPU-GPU integration

500 million graphic units sold in 2011

- 65 million graphics cards (5 million high end)
  15 B$ revenues

AMD and Nvidia use both the same foundry for the GPU production (TSMC).
28nm structure size production has started → some production and yield problems

High end cards for gaming, engineering and computing → niche market
Discrete graphics market under pressure from INTEL (Sandy-Bridge, Ivy-Bridge)

Focus is on energy efficiency and the mobile market
→ Tegra and Fusion
→ New Kepler from Nvidia with ‚reduced' gaming and DP performance

New gaming model based on ‚cloud' computing (OnLive Gaikai)
High performance gaming on mobile devices: less graphics cards needed

Discrete GPU costs will rather stabilize than heavily decline

Tape Storage

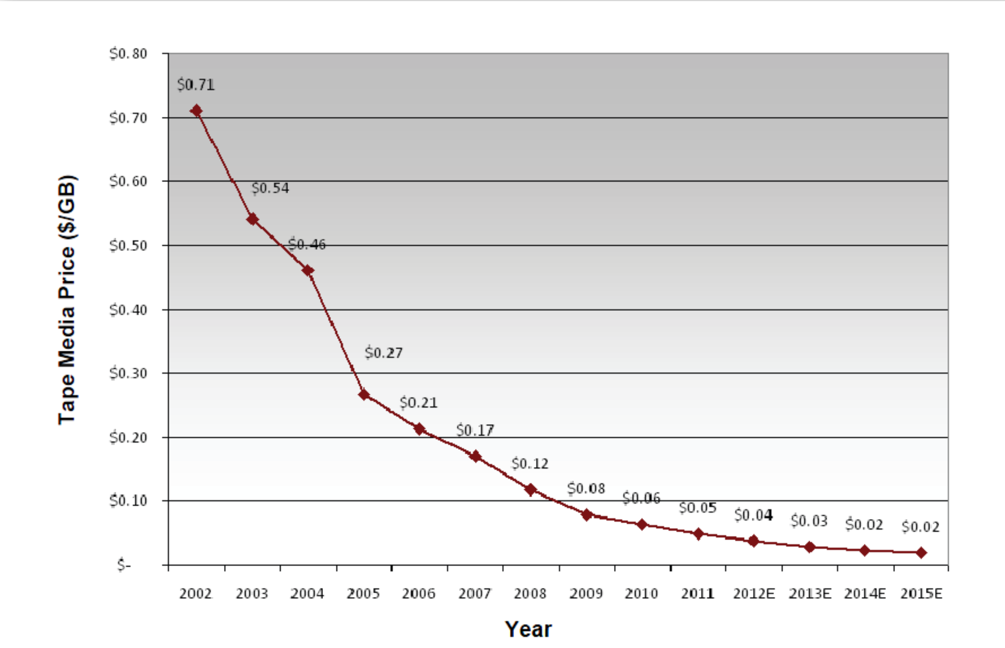

Tape market : 3.5 B$ revenue per year

Today 3 main tape formats:

LTO	LTO-5	1.5 TB
IBM	3592JC	4.0 TB
Oracle	T10000T2	5.0 TB

LTO covers 90% of the market share

Current media cost: 0.03-0.04 Euro/GB

Solid technology roadmaps exists for all 3 competitors
→ 10-15 TB cartridges in 2017

Large scale backup

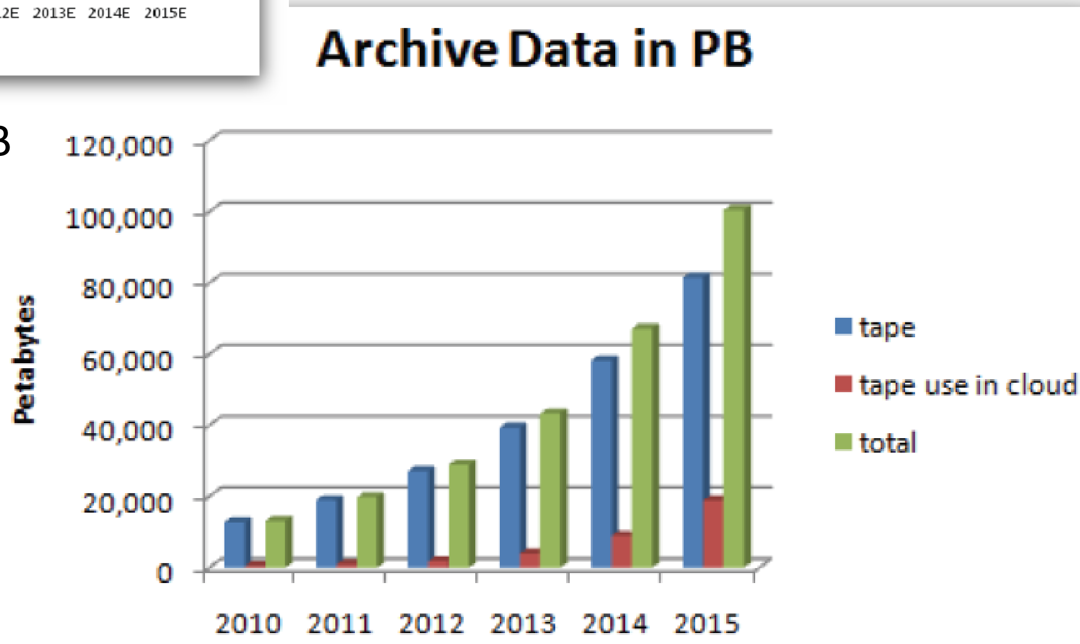

General Trends

Cloud Computing
Cloud Storage

Large scale storage and data access
Large scale processing, data mining and data analysis
Fast feedback

Mobile Devices:
Smartphones, Tablets

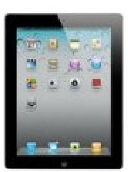

Low data volume
Low processing capacity

Internet of things:

Monitoring Sensors and actuators
surveillance, security, healthcare, biosensors, transport, environment, Geolocation, …….
Citizen Science, do-it-yourself
→ Tinkerforge, Arduino, Ninja Blocks,…

"More than Moore"
Integration of digital, analogue and micro-mechanical functionality

‘Exotic’ Computing Technologies

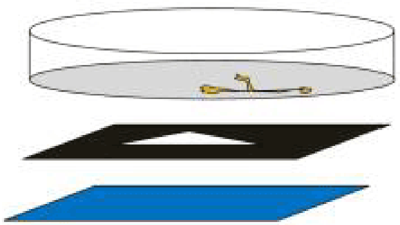

- Implementation of AND/NOR gatters based on the movement of soldier crabs

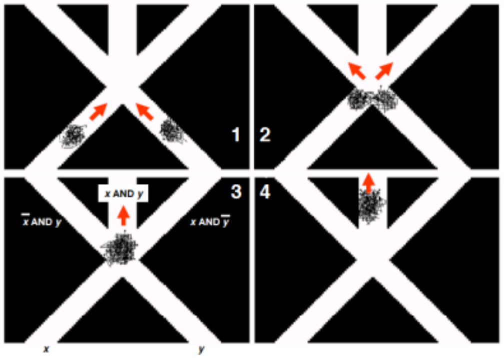

- Biocomputing using single-cell organism (plasmodium) e.g. Shortest way through a maze
- Chemical computing with traveling wavefronts of crystallization in supercooled sodium acetate
- Chemical oscillating patterns based on the Belousov-Zhabotinsky reaction
- Molecular Cryptosystem for Images by DNA Computing

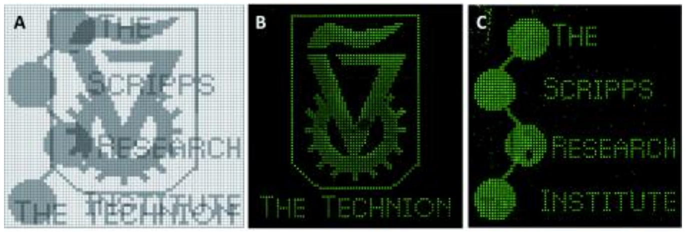

- Some current breakthroughs is quantum computing: reaching 14 qbits, 100 microseconds coherence times, moving from atoms to nano-dots, but still 10 years away from a usefull production system

Summary

Moore's Law is still alive
(But Moore's Law and server costs have a complicated 'relationship')

Technology roadmaps are well advanced, approaching 10nm structure sizes will be challenging and very expensive (~2016)

Constant flow of very exciting bleeding edge research results, but no market relevance

Only very few companies are sharing the various markets (CPU, Disk, Memory)

The market is pushing for mobile device and energy efficiency

Continuing price/performance improvements over the next years, but indications for a change in the slope (less steep)

Prediction is very difficult, especially about the future.
Niels Bohr