# Technology, Market and Cost Trends 2012

Bernd Panzer-Steindel, CTO CERN/IT
panzer@mail.cern.ch
May 2012

## Table of Contents



# 1. Introduction

The following paper tries to give an overview of technology and market trends in computing. The final goal of this exercise is to provide essential input for the budget and resource planning of the CERN computer centre. 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.

- 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

The future evolution of the architecture and amount of installed resources in the CERN computer centre depends on the requirements from the LHC experiments, the market trends affecting the cost development of computing equipment, the available budget and the infrastructure boundary conditions (buildings, purchasing  conditions, power and cooling, politics, software inheritance, etc.).

The derived cost predictions for the various computing parts have of course been made with today's state of information and thus have to be judged with some care.

*"There is no reason anyone would want a computer in their home."*
Ken Olsen, Founder of DEC, 1977

"***Apple is already dead.***"
Nathan Myhrvold, former Microsoft CTO, 1997

"***Prediction is very difficult, especially about the future.***"
Niels Bohr, Danish physicist (1885 - 1962)

The High Energy Physics data processing and data storage is done with low-end server nodes. The basic features of the data processing are: program performance is determined by integer calculations (80%); physics events are independent, thus no fine grain parallelism is needed; processing programs need >= 2GB of memory per job (=core).
The CERN computer centre provides roughly 15% of the needed High Energy Physics processing and data storage capacity worldwide, while 85% of the capacity is provided by several hundred sites as part of the worldwide_LHC-GRID.
Here is a brief overview of the installed equipment in the CERN computer centre:

- 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, 65 PByte data on tape
- 3.5 MW power and cooling
- Replacement and purchase rate is about 1500 servers per year (in AND out)

# 2. Market Trends

The revenues in the worldwide semiconductor (CPU, GPU, memory, flash, microcontroller, etc.) market reached 311 billion $ in 2011. Analysts expect in the coming years moderate growth rates of 2-4%. Figure 1 shows the detailed list of the largest companies in this market **[Link 1].** INTEL is in terms of revenues the biggest company and only 10 companies are sharing 50% of the worldwide revenues. This kind of consolidation is a common trend.

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

***Figure 1: Worldwide Revenue Ranking for the TOP-25 Semiconductor Suppliers in 2011***

The computing market can be separated into three different categories: Servers, Personal Computer and Mobile Devices.
The following table describes the detailed market characteristics of these categories and sub-categories for the year 2011. These numbers are regularly published by the main analyst groups (IDC, Gartner, iSupply, etc.), but they sometimes vary by 5-10% from company to company.





| Type | #units sold [million] | Revenues [billion $] | Growth rates [ % ] |
|---|---|---|---|
| | | | |
| 1. Servers | 9 | 52 | +8 |
| HPC server | | 10 | +8 |
| Supercomputer | | 4.4 | +8.5 |
| | | | |
| 2. Personal Computer | 351 | 260 | |
| Desktop | 112 | | +2.3 |
| Notebook | 210 | | +7.5 |
| Ultrabook | 4 | | |
| Netbook | 29 | | -25 |
| | | | |
| 3. Mobile Devices | | | |
| Smartphones | 490 | 160 | +63 |
| Tablets | 63 | 26 | +274 |
| | | | |

***Table 1: Computing market details (categories, revenues, growth rates)***

The clear trend to mobile devices can also be seen by the following diagram from IDC, predicting the unit shipments of so called Smart Connected Devices. **[Link 2]**

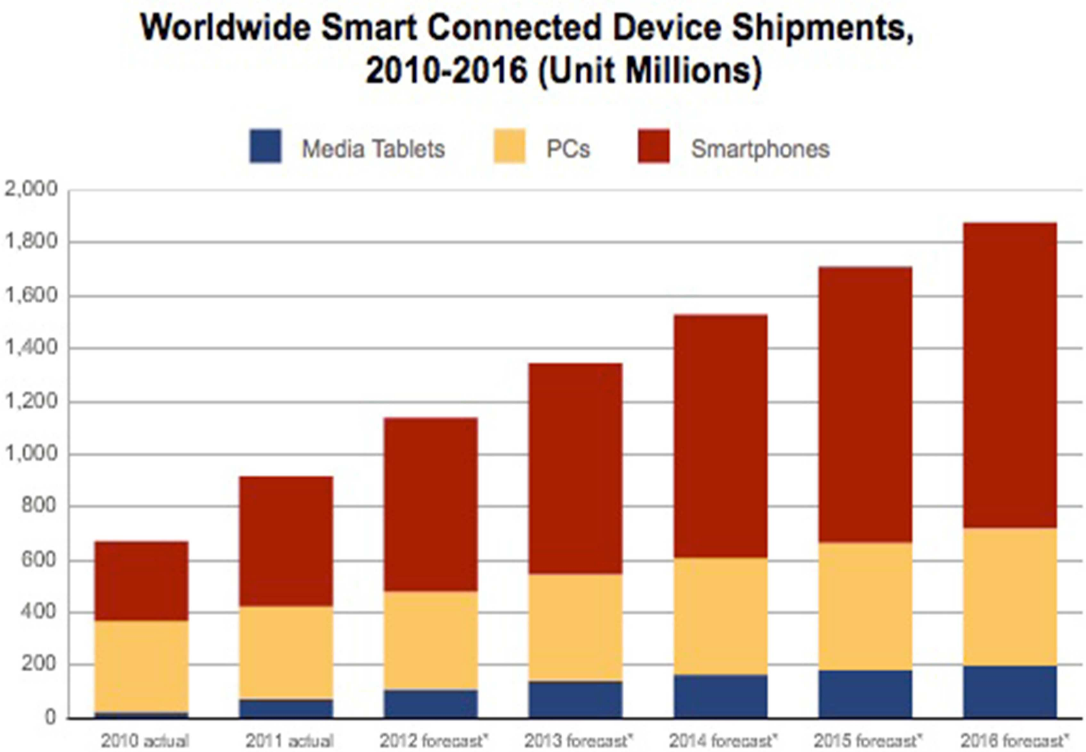


***Figure 2: Worldwide Smart Connected Device Shipments, 2010-2016***

In April 2002 the one billionth PC was shipped to customers, the second billion was reached in 2007 and by this time (May 2012) probably 3.5 billion PCs were manufactured and shipped worldwide. But this translates to only about 1.5 billion PCs installed today and an expected 2 billion in 2015.
This has to be compared with the existing ~6 billion mobile phone subscriptions at the end of 2011.

There is a clear market push from raw performance to power efficiency (execution and standby): few cores, simple processors, specialized processors (DSP, FPGA), SSD disks, NAND flash memory,……
It is sort of ironic how much ingenuity and effort is used to produce more and more complex silicon chips and at the same time a similar amount of sophistication in software tries to minimize the usage of that silicon.
This "Dark Silicon" phenomenon **[Link 3]** has site-effects on the possible energy efficiency and use of multi-cores on silicon chips.

A few other observations:

- The number of sold Smartphones is now higher than the amount of PCs sold
- Smartphones and tablets have the highest growth rates
  → this 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 fewer units were shipped.
  → Lucrative niche market for HPC and Supercomputer
  → Consolidation and efficiency improvements (virtualization) in large computer centres is affecting the server market

# 3. Processors

In this chapter about server processors I will concentrate on INTEL processors and more or less ignore AMD products:

- The INTEL server market share is ~94% compare to ~6% for AMD.
- There are also products from IBM (PowerPC), Oracle and Fujitsu, but aimed at the HPC niche market
- INTEL has a 12-18 month technology lead in processor technology
- There are strong indications that AMD will 'leave' the server market, but maybe still be present in the HPC area

**[Link 4]**

The complexity increase of processors (CPU and GPU) is still following Moore's Law, i.e. number of transistors on a chip is doubling about every 18 month.

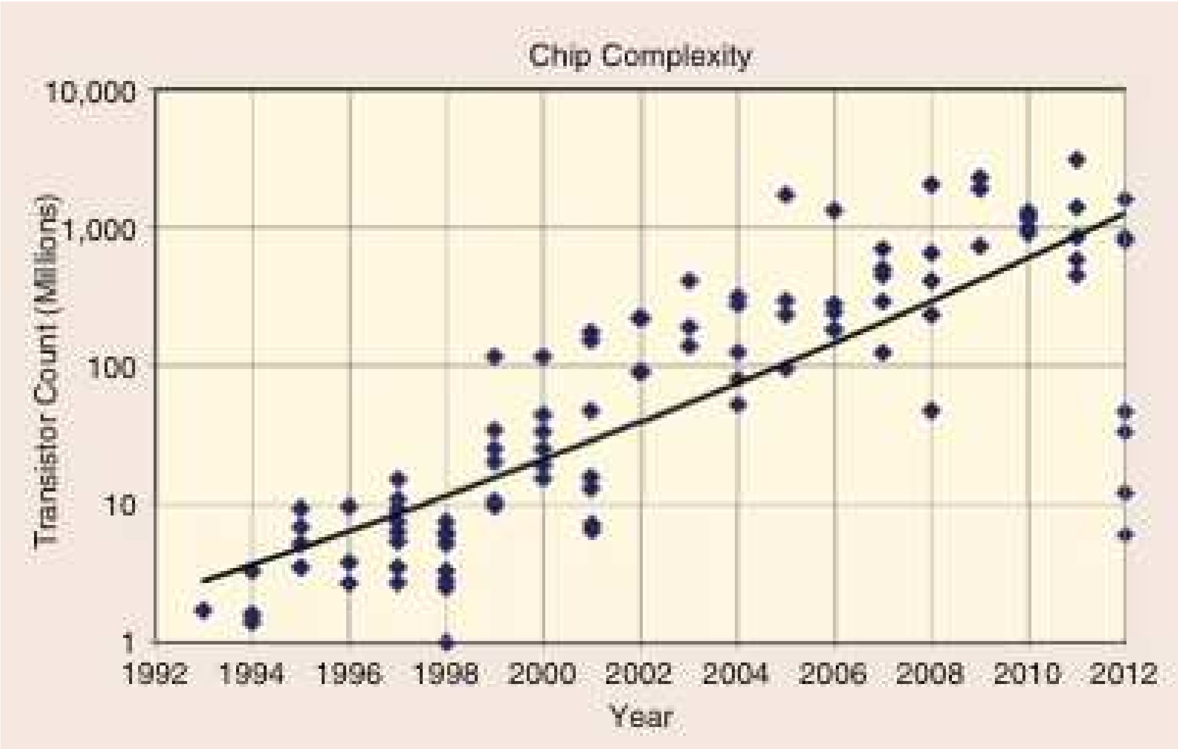


***Figure 3: Chip complexity evolution, 1993 - 2012***

The diagram from the 2012 International Solid-State Circuits Conference (ISSCC ) trend report shows the chip complexity evolution during the last 20 years. [Link 5]
The most complex processors are coming from the current generation of Graphics Processing Units : Nvidia 680 GTX with 3.5 billion transistors and AMD Radeon 7970 with 4.3 billion transistors.
This list just covers processors, if one would include chips is general, than the Virtex-7 FPGA from Xilinx would be the 'winner' with 6.8 billion transistors.

Today all foundries are using wafer sizes of 300mm (12 inch) in their production lines. The full transition to 450mm was planned for 2012, but has now moved to 2017-2018. The estimated cost of a new 450mm fab manufacturing 11nm transistors is on the order of $US 10B, putting the cost out of reach to all but a select few manufacturers (3-4). [Link 6]

The GPU units from Nvidia and AMD are currently the most complex processors. Both chips are using 28nm structure sizes for the transistors and memory units. The PC processors are fabricated in 32 nm (AMD) and 22 nm (INTEL) technology and also the latest ARM processors are manufactured with 32nm (e.g. Samsung Exynos) transistors. The manufacturing of chips with structure sizes of 28-32 nm belongs to the same technology generation (full-pitch and half-pitch); the same is true for the structure range of 19-22 nm.

The focus of the manufacturing is the reduction of leakage currents and running the chip with as low voltages as possible. Each step to shrink the structure sizes on a chip had to be accompanied by sophisticated advances in material science:

- 90 nm → strained silicon
- 45 nm → high –k metal gates (hafnium oxide)
- 22 nm → 3D tri-gate transistors

The 22nm process with FinFET transistors includes the usage of strained silicon and high-k metal gates.
[Link 7]
These structures are still created with the optical immersion-lithography process based on 193 nm light sources.

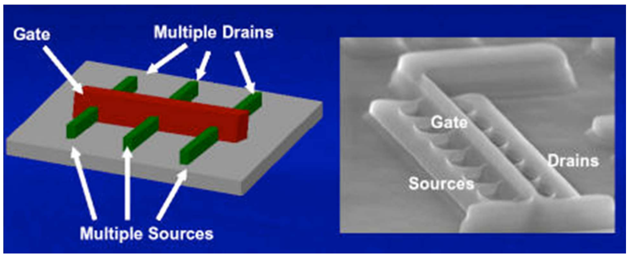


***Figure 4: 3D tri-gate transistors with 22nm structure size***

The next step for the processor manufacturing technology is the move to 14nm structures on the chip. Figure 5 shows a leaked roadmap for the Intel server microarchitecture crossing the 10nm boundary some time in 2018. [Link 8]
This is confirmed with some more official announcements. [Link 9]

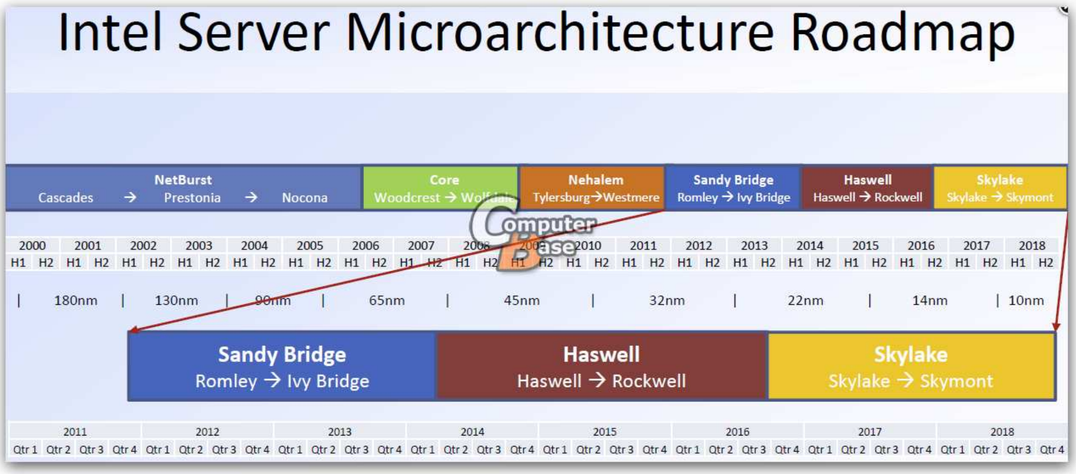


*Figure 5: Leaked Intel Server Microarchitecture Roadmap 2011*

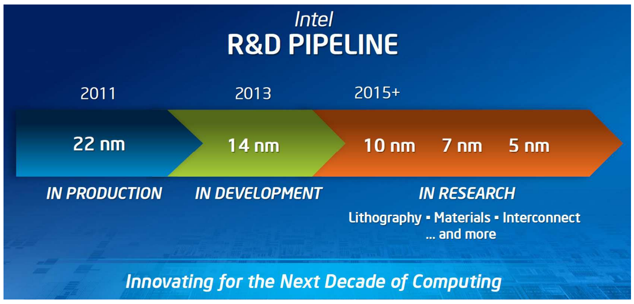


*Figure 6: Intel Structure Size Roadmap from 2012*

It is very difficult and expansive to create the next shrink of structure sizes (<14nm) as the technology has to move to EUV (extreme Ultra-violet) lithography. This is still under development and the current target for production in 2015 will be hard to achieve. [Link 10]

## 3.1. Multi-Core Evolution

It is important to understand the evolution of the number of cores per processors and CPU server, because there are several side effects:

- The number of IO streams is directly correlated to the number of cores (in general one data processing task per core). Thus the IO performance of the local server storage needs to be matched (number of spindles, SSD versus HDD, costs)
- The overall processing and storage architecture in the computer centre needs to map the total number of IO streams (cores) with the capability of the storage system (number of spindles, SSD caches, network, etc.)

- Due to software requirements the processing core needs to be equipped with a certain amount of memory

Diagram 7 shows the evolution of the core-count per processor for three different categories:

1. The core count of processors used in the servers installed in the CERN computer centre over the last 7 years
2. The core count of INTEL processors in the DP (dual socket) Xeon series using the official release date of the processor
3. The core count of INTEL processors in the MP (high end, >= 4 socket) Xeon series using the official release date of the processor

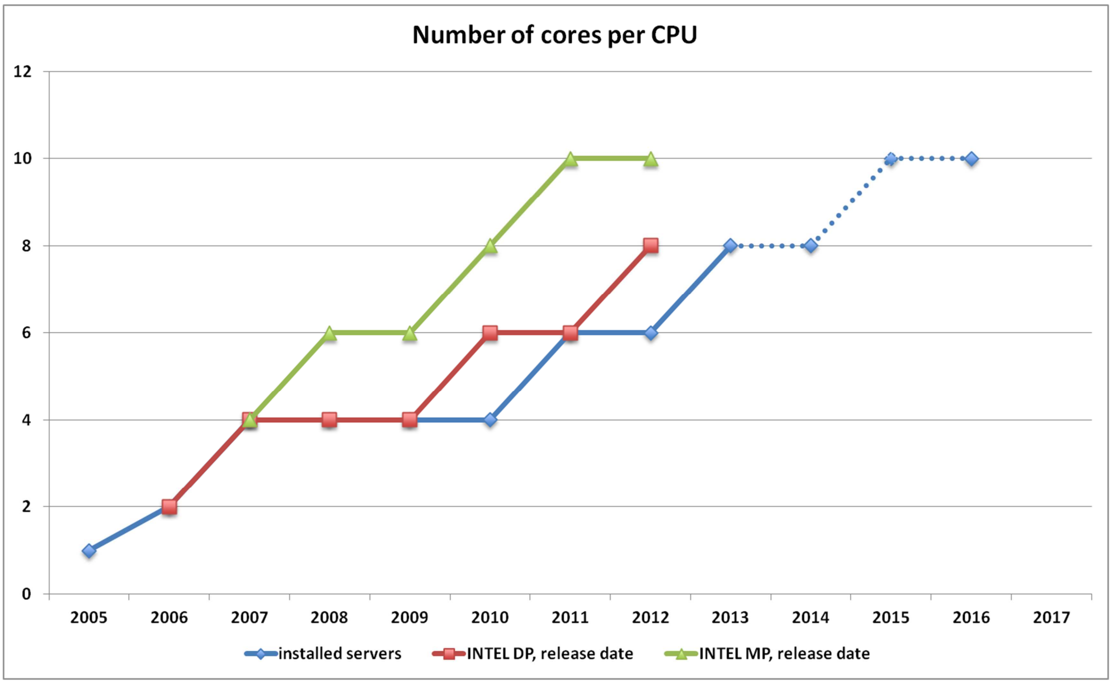


***Figure 7: Evolution of the number of cores per CPU (INTEL)***

The reference is always physical cores and not 'logical' cores like in the SMT case of Intel processors. This would double the amount of cores but provides only about 25% more performance. AMD has in general twice the amount of cores than INTEL as they can reach only similar performance values per processors by having more core running at higher frequency. The Interlagos AMD processor has for example 8*2 cores which is actually somewhere in-between full physical cores and the INEL SMT cores.
Figure 8 shows an extract of the 2012 ISSCC trend report report [Link 5] confirming the observation that the core count evolution is following a more linear than exponential growth rate of about one-two cores per year. The high core values in 2012 include special chips like the Tilera 64-core many-core system.

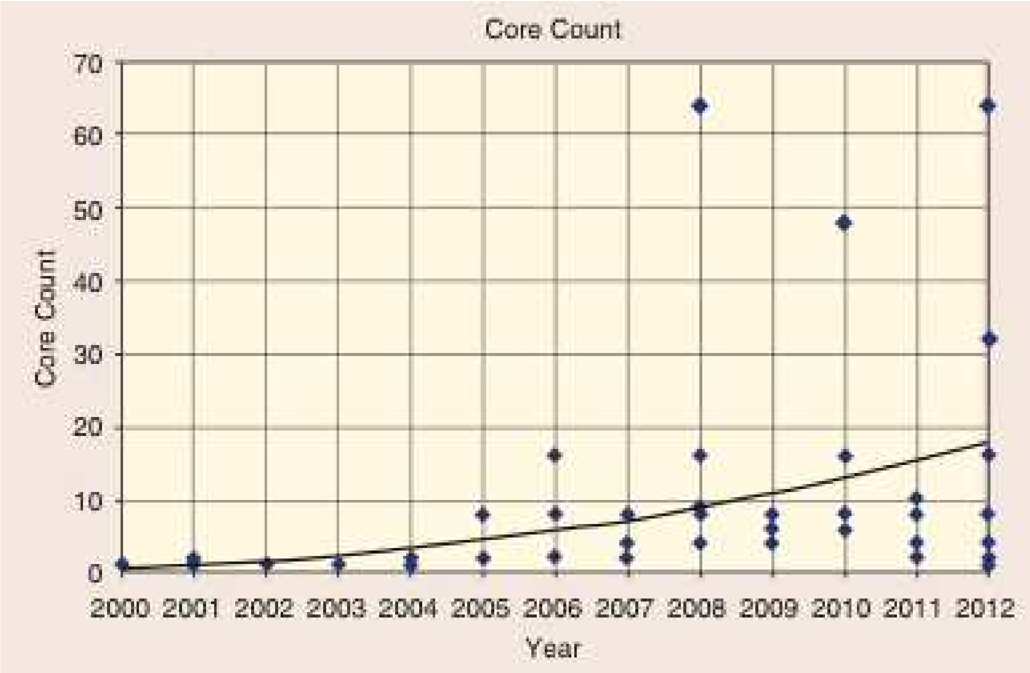


*Figure 8: Evolution of cores counts per CPU, 2000-2012*

The market penetration of multi-core systems shows a similar trend. In 2011 only 60 million out of the 490 million sold smart-phones were two core systems and less than 30% of the sold notebooks had a quad-core processor.

The following article provides some more details about the multi-core problems the industry is facing: [Link 11]

## 3.2 Graphics Processors

The market for graphic processors reached 15 billion $ revenues in 2011.
There are three different types of graphics cards: discrete PCIe graphics cards, on-board (motherboard) GPUs and integrated CPU-GPU combinations.
In 2011 a total of 500 million graphic processors were sold, of which 65 million were on discrete PCIe cards.
There are today only three companies in the GPU market: INTEL, AMD and Nvidia.


The high-end graphics cards are becoming more and more popular in the HPC computing area due to their extremely high floating point processing performance. There are 35 supercomputer using Nvidia GPUs in the top500 list from Nov 2011. AMD has very performing graphics chips, but the Nvidia software eco-system around CUDA is much better organized and has attracted a much larger and active user community. Actually the

AMD cards are providing today a much better DP versus SP ratio than Nvidia (Single and Double Precision floating point operations).

The latest graphics card from Nvidia called Kepler has 1536 cores and has a single-precisions floating point performance of 3090 GFlops. [Link 12]
Like the processors from AMD it is manufactured by TSMC using 28nm structures. In addition to the Nvidia 680 GTX aimed at the gaming market, new designs with better DP (double precision) floating point performance have been announced for the second part of 2012.

Table 2 shows a comparison between the latest high end graphics cards and an Intel Xeon server processor.

| | #cores | Frequency [GHz] | #transistors [billion] | GFlops single | GFlops double | Structure size [nm] |
|---|---|---|---|---|---|---|
| **Nvidia Tesla 2090** | 512 | | | 1331 | 665 | |
| **Nvidia 680 GTX** | 1536 | 3 | 3.54 | 3090 | 129 | 28 |
| **AMD Radeon 7970** | 2048 | 2.75 | 4.3 | 3788 | 947 | 28 |
| **Intel Sandy Bridge** | 8 | 2.6 | 2.27 | | 152 | 32 |
| | | | | | | |

***Table 2: Floating point performance comparison between Nvidia, AMD and INTEL GPUs/CPUs***

The interesting observation is the low double precision floating point performance of the 680GTX. Nvidia is starting to diversify their product pallet and creating extra lines aimed explicitly at the HPC market. The mentioned first Kepler product (680 GTX) is meant for the high end gaming market (high graphics performance == single precision floating point and low power consumption), but featuring a low DP performance.
Both companies are focusing their activities on the mobile market and the corresponding energy efficient graphics processing: Tegra from Nvidia and Fusion from AMD.
High-end graphics cards for extreme gaming, engineering and HPC computing are a niche market with only ~5 million units sold in 2011.
The major market for discrete PCIe graphics cards is the 19 billion $ PC games market. But this market is getting under pressure from several developments:

- The integrated CPU-GPU processor combinations from INTEL (e.g. Ivy Bridge) and AMD (APU concept) are providing better and better graphics performance at higher energy efficiency and lower cost
- New methods to provide high performance gaming on mobile devices, especially tablets, via cloud computing (Onlive, Gaikai, Nvidia cloud streaming) . This still requires high end graphics cards at the cloud provider, but this can be organized in a much more efficient central manner, reducing overall the required amount of graphics cards

[Link 13]

Thus high-end graphics cards will become most likely even more of a niche market with a less optimal price/performance tag.

## 3.3 New CPU servers and Many-Core systems

There are quite a few activities in the area of new processor and new server designs. The following list tries to give a selected overview.

**Tilera:** 64bit processors in a two dimensional array; multiple mesh networks; just released a 36-core 64bit processor version; they made some successful tests with Facebook applications and achieved a factor 3 better transactions per power unit, but there seems to be no deployment plan. [Link 14]

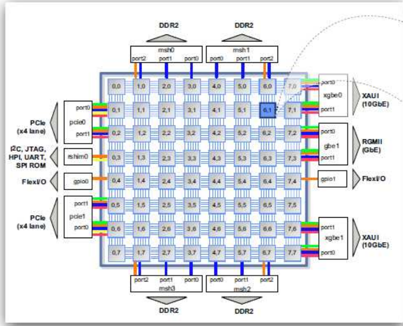


**Adapteva:** Epiphany IV processor; 64 RISC cores, aimed at special high performance applications, like speech recognition in mobile devices. [Link 15]

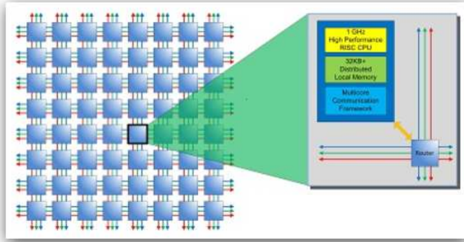


**ZiiLabs formerly known as 3DLabs**: 100-core system; ARM Cortex-A9 + 96 StemCell Media Processing cores. [Link 16]

**Fudan University**: 16-core RISC processor using message passing. [Link 17]

**Plurality:** 256-cores Hypercore processor; used for wireless infrastructure Company founded in 2004; first processors in 2008; not much development since then; last press releases in 2010. [Link 18]

**Neuromorphic processors** e.g. silicon retina ‚Third Eye' motion detection, SyNAPSE project using memristors; combining sensors, analogue cicuits and processors. [Link 19] [Link 20]

**Kalray**: European startup company (Orsay France); 256 processor array on a 28nm CMOS chip; no activity since quite a few month. [Link 21]

**INTEL:** Many-Integrated-Cores (MIC) Knights Corner; 50+ cores; will become a product in 2013; focus is the HPC market; co-processor. [Link 22]

**Enterpoint:** PCIe co-processor board based on Xilinx FPGAs; 6 processor cells [Link 23]

**Maxeler supercomputer**: In production by JP Morgan for high frequency trading using Intel Xeon processors plus 4 Xilinx Virtex-6 co-processor cards. [Link 24]

**BittWare**: PCIe board using Altera FPGAs, used for high frequency trading in the financial acceleration market. [Link 25]

In general there have been quite successful tests with FPGA systems in the HPC area. The achieved price/performance and actual absolute performance was factors higher than provide by standard x86 processor systems. The main obstacle is the very specific and complicated FPGA programming model and limited and not very user-friendly programming tools. [Link 26]

**Calxeda:** Servers based on ARM Cortex-A9 (SoC);
used by HP (Project Moonshot); 288 processors in 4U. [Link 27] [Link 28]

There are activities to incorporate the low power ARM processors in the HPC processing environment. [Link 29]

**SeaMicro**: Micro-server design originally using low power Intel processors, special motherboards and a proprietary interconnect; the company was lately bought by AMD and will now provide servers with low power AMD processors. [Link 30]

There are obviously quite a few companies investing into many-core products and micro-servers based on low power x86 processors. On one side these new systems are tailored to specific clearly defined areas (audio processing , video processing, network, encryption, etc.) or have to deal with the usual software problems of highly parallel computing systems. The latter problem is of course well known and described [Link 31] [Link 32] and there are specific efforts like the IBM Renaissance Project [Link 33] to improve the situation.

Some of the earlier many-core projects were less successful and there are today still large discrepancies between the announcement of high many-core processors and there introduction into the market.

- The Intel Larrabee 80 core processor announced in 2007 and canceled in 2009
- The Cell processor; a collaboration between Sony and IBM and first used in the PS3 game console in 2006; IBM stopped the development in 2009
- The Intel 48-core SCC cloud computer; announced in 2009, but no product yet

# 4. Storage overview

The flowing table shows key performance figures for the 5 different storage technologies and markets in 2011.

| Technology type | Revenues [ billion $] | #units shipped [million] | Sold storage size [ExaBytes] |
|---|---|---|---|
| | | | |
| **DRAM Memory** | **31** | **800** | **2** |
| | **Samsung, Hynix, Micron 91% market share** | | |
| **NAND Memory** | **30** | **4000** | **20** |
| | **Samsung, Toshiba, Micron, Hynix 99% market share** | | |
| **Solid State Disks** | **5** | **17** | **3** |
| | **> 50 companies** | | |
| **Hard-Disk-Drive** | **28** | **630** | **350** |
| | **Western Digital 37%, Seagate 47% , Toshiba 16%** | | |
| **Magnetic Tape** | **1** | **27** | **20** |
| | **LTO-Consortium, IBM, Oracle** | | |

***Table 3: Storage market overview (companies, revenues, shipments)***

All storage device markets (besides SSDs) are in the hands of only 3-4 different companies. The last years have seen a strong consolidation of these markets.

The following three figures give some overview of the technology evolution of these storage products. The shrinkage of the feature sizes and the corresponding areal density increase is still following exponential curves, but it is expected that these high percentage growth rates will decline in the future. This is due to the technology problems when feature sizes are approaching the 10nm range. In addition large-scale investments are needed to build the next generation fabrication units. These very interesting predictions and plots are from an IBM presentation by R. Fontana and colleagues [Link 34].

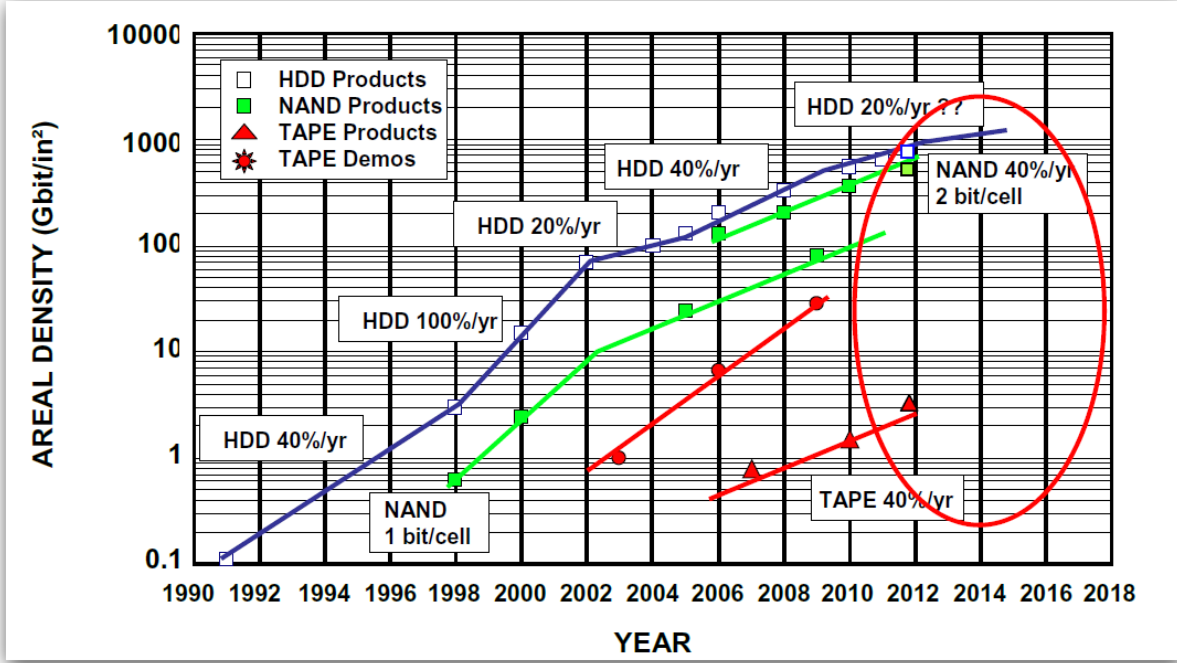


*Figure 9: Areal density evolution of HDD, NAND and TAPE*

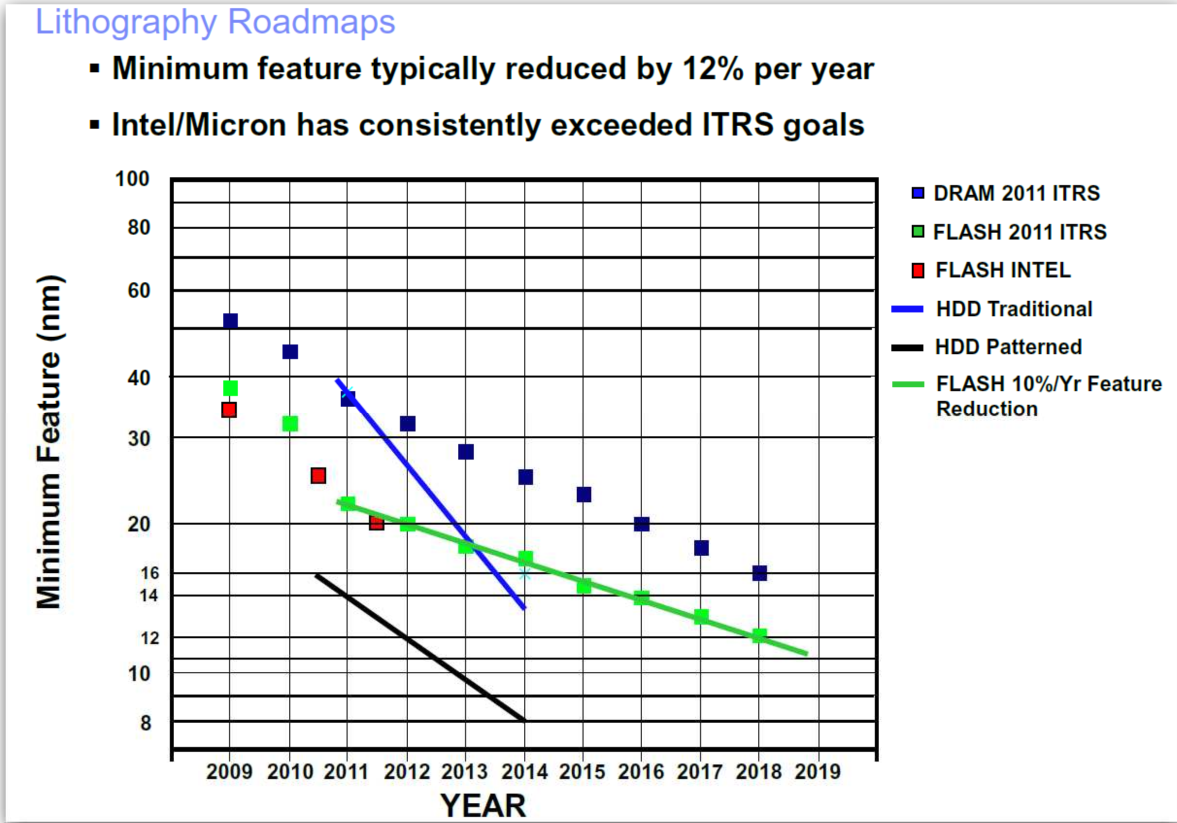


*Figure 10: Feature size evolution for HDD, DRAM and NAND memory*

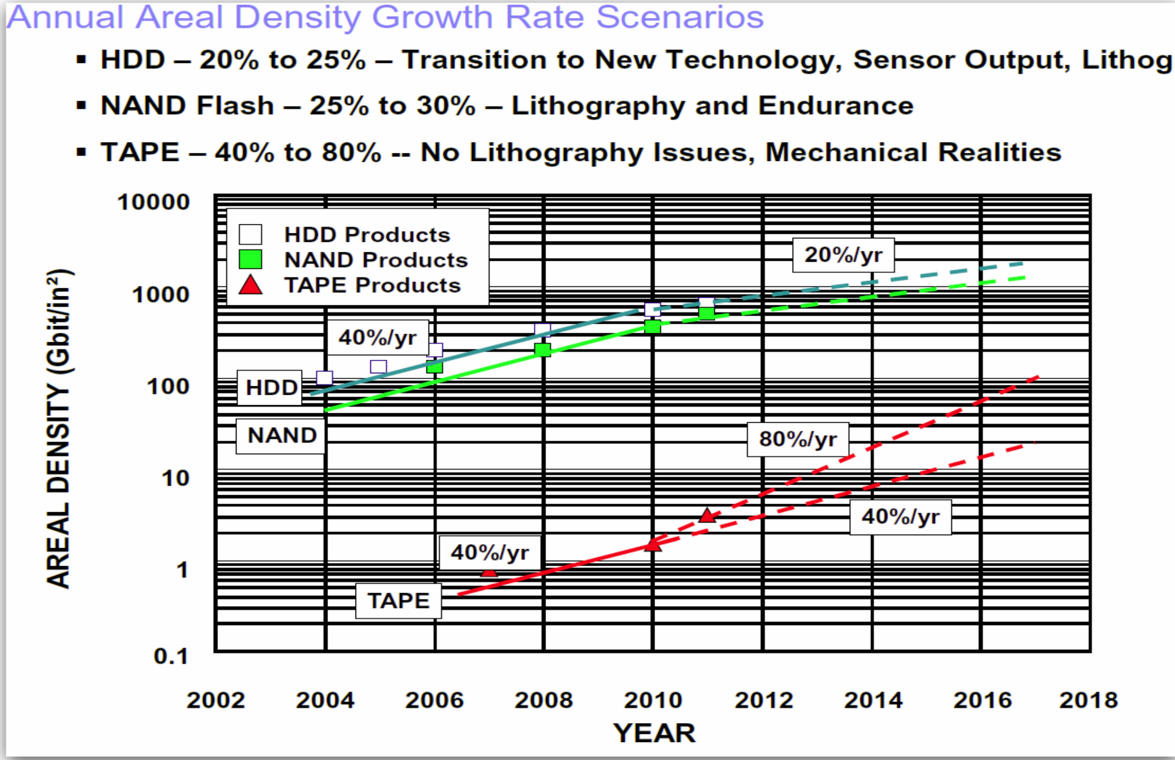


*Figure 11: Annual areal density growth rate scenarios for HDD, NAND and Tape*

## 4.1 Memory

The revenues in the DRAM memory market were at the 31 billion $ level in 2011, but overall there was a decrease of 24% in revenues. There were several reasons for this decrease: overproduction, low growth rate in the PC market and high growth rates in the mobile market, which uses NAND flash.
There are only three companies sharing the market (Samsung, Hynix and Micron). The price evolution of memory modules is very volatile and can have large fluctuations due to the dependencies on other market factors.
Figure 12 shows the street price evolution of server quality memory units (ECC, registered) for the last 6 years.

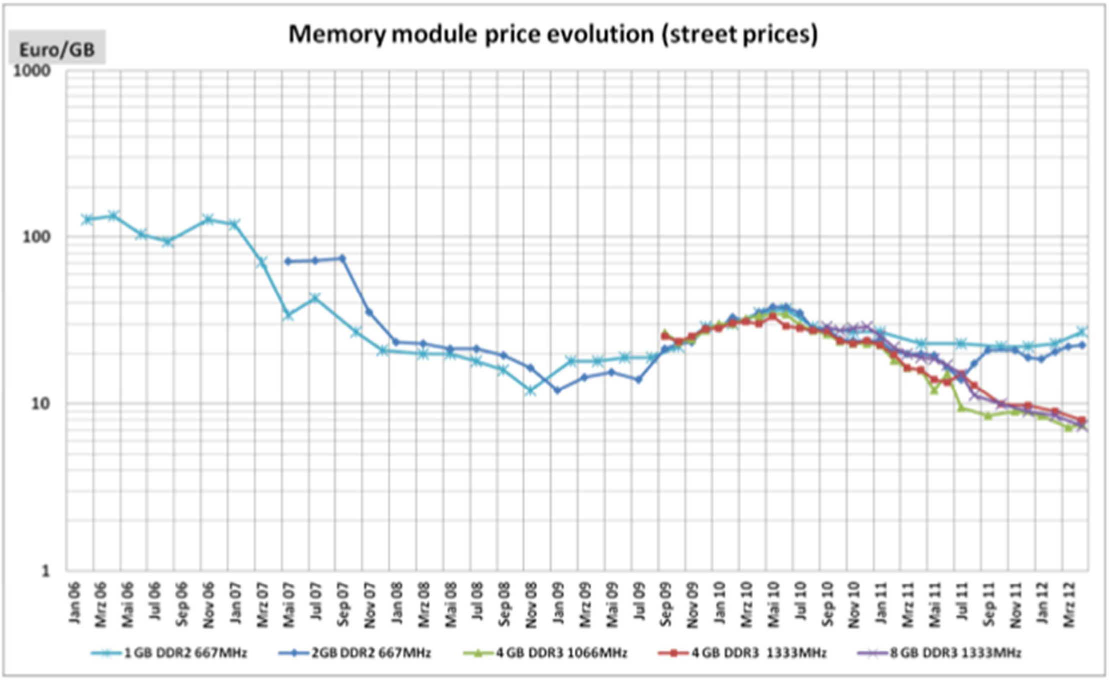


***Figure 12: Street price evolution of server DRAM memory modules***

The dominating DRAM type in today's serves is DDR3, which is manufactured using 30-40 nm processing technology. There is a clearly defined roadmap for the performance and technology evolution of DRAM. These roadmaps are regularly published by JEDEC (Joint Electron Device Engineering Council), which is developing open standards fro the microelectronic industry.
Table 4 shows the latest forecasts [Link 35] [Link 36]:

## 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.

***Table 4: DRAM memory roadmap***

The move to DDR4 (1866 MHz) memory will focus on clients in the high-end server market. The HPC market has more and more a problem of matching the processors speeds with the memory throughput (memory wall problem) [Link 37]. This is not so relevant for high energy physics applications, as they have only limited needs for memory bandwidth.

The future technology roadmaps of manufacturing RAM include the new 3D transistor technology [Link 38] and the implementation of hyper memory cubes [Link 39].

## 4.2 NAND flash memory

Global NAND flash revenues are projected to reach US$22.9 billion in 2012, up from US$21.2 billion in 2011. And with NAND flash consumption increasing in the three principal markets for smartphones, tablets and SSDs in Ultrabooks, NAND revenues will climb continually during the next few years, hitting approximately US$30.9 billion by 2016, IHS said. [Link 40]

Smartphone shipments this year will hit 626 million units, with an average NAND flash content of 9.0 gigabytes per unit.
On its own, Apple will consume about 25% of the overall NAND supply in 2012, equivalent to some eight billion gigabytes, IHS predicted. The iPad will be responsible for 74% of NAND consumption in the tablet segment this year, accounting for 2.8 billion gigabytes out of a total 3.8 billion gigabytes in tablets.

The NAND memory technology comes in three flavours: SLC (single-level cell, 1bit), MLC (multi-level cell, 2bit), TLC (three-level cell, 3bit). The manufacturing process has now reached the sub 20nm area and mass production of these NAND chips is already expected in 2012. Toshiba and SanDisk lately announced their new 128 Gbit flash chip using 19nm structure sizes and three-bit cells (TLC). [Link 41]
The only way to decrease the cost of NAND memory storage is the shrinkage of the structure sizes and the increase of bits per cell. But this causes severe endurance problems, the possible P/E (program-Erase) cycles per cell decreases rapidly as can be seen in figure 13. [Link 42]
The further shrinkage of the structure sizes have severe consequences for the error rates and write latencies. A recent research paper explains the clear limitation of this storage technology [Link 43].

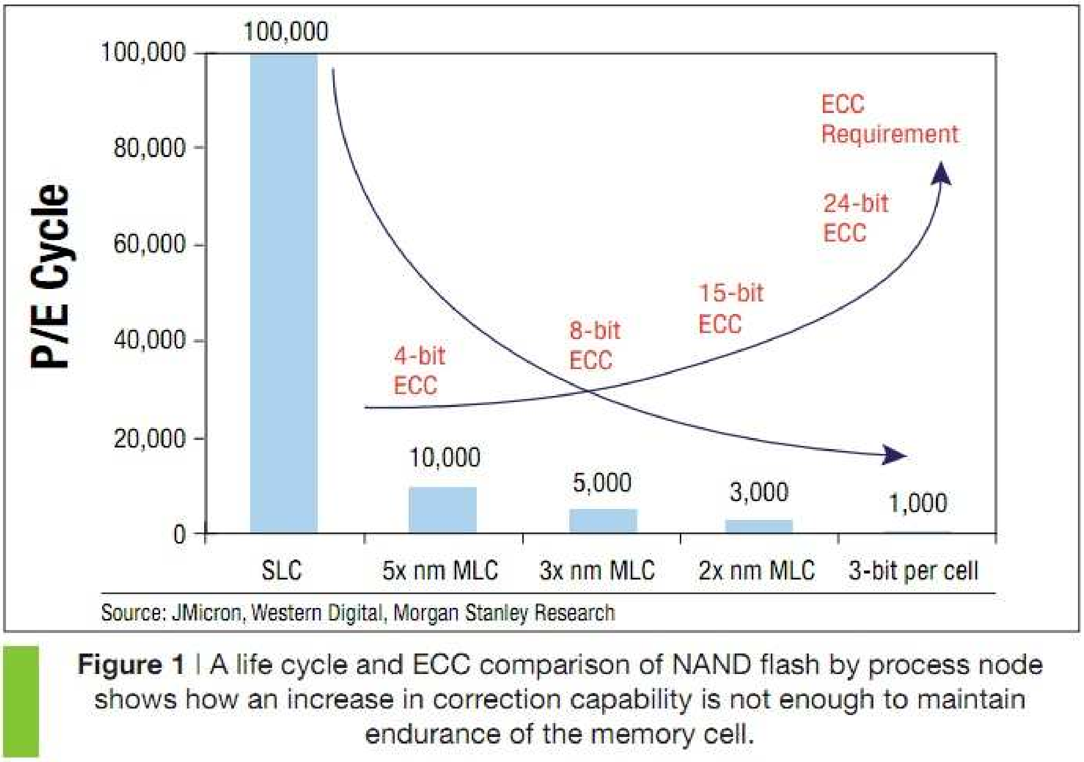


***Figure 13: Decrease of NAND flash memory endurance with shrinking feature sizes***

According to new research from International Data Corp. (IDC), the worldwide solid state storage industry revenue reached $5 billion in 2011, a 105% increase from the $2.4 billion in 2010 and IDC expects the market will expand further in 2012 and beyond. IDC expects worldwide SSD shipments to increase at a compound annual growth rate (CAGR) of 51.5% from 2010 to 2015.
SSDs will account for some 3.3 billion gigabytes of NAND flash consumption this year, up from 1.7 billion gigabytes in 2011, IHS said. [Link 44]

The cost for an average Solid State Disk (MLC consumer level) has dropped considerably during the last 2 years and one can find SDDs as low as 1$/Gbyte. But there are still about 50 companies offering various versions of SSDs with a cost bandwidth of at least a factor 20, with enterprise SLC disks being the most expensive type. The endurance limitations especially of TLC type NAND memory has to be compensated by sophisticated wear-level algorithms in the SSD controller (high bit ECC, compression, >100% spare space, delayed write, etc.). Thus one finds today only 3 main brands of controller types in the market: Marvell, Samsung and LSI (Sandforce).

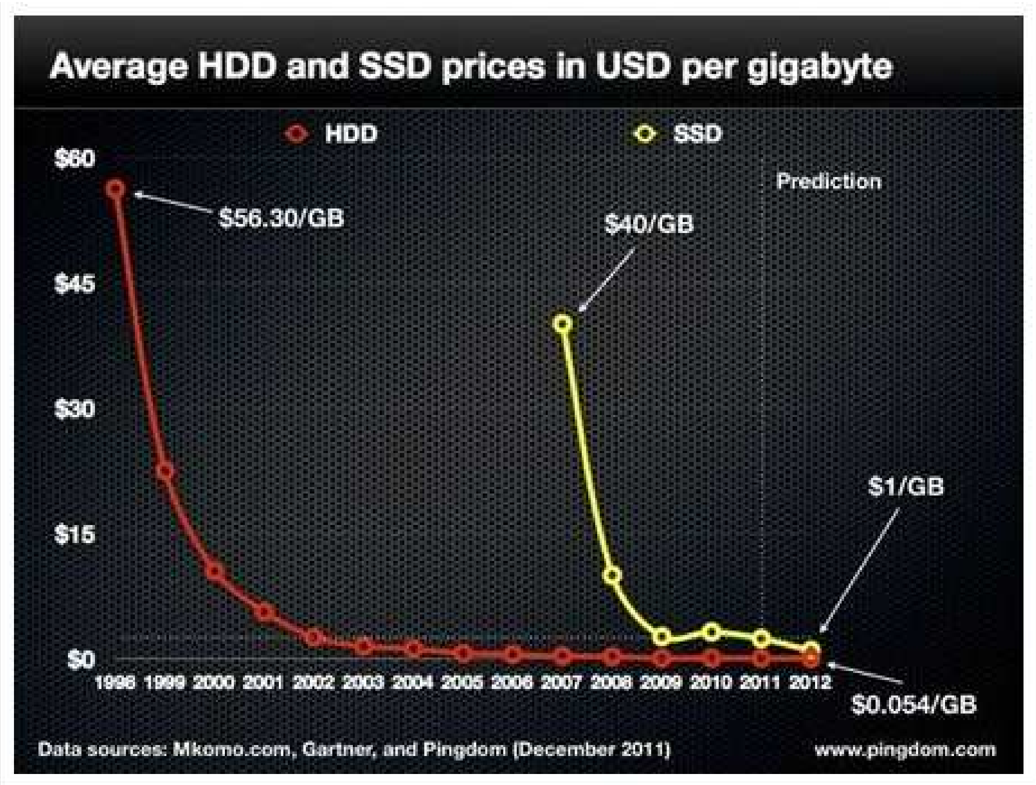


***Figure 14: Price evolution of HDD and SDD space***

The endurance problems are posing a considerable limitation for the near-term future and it is expected that SDD disks will never approach the same price/Gbyte levels as HDDs. This can only happen when the new technologies like PCRAM or memristors will reach the mass market, which will still take many years.

## 4.3 Hard Disk Storage

During the last 2 years there has been some considerable HDD market consolidation. [Link 45]
Only three companies are now sharing the ~30 billion $ HDD market. Due to the Thailand disaster the average growth rate in 2011 was -4.5%. The forecast for 2012 and onwards is expecting positive growth rates of 7-8 %. The major demand will come from cloud storage, which today is already at the 50% level. The estimated demand for HDD storage was about 400 Exabyte in 2011 and should reach about 1 ZettaByte per year in 2016.

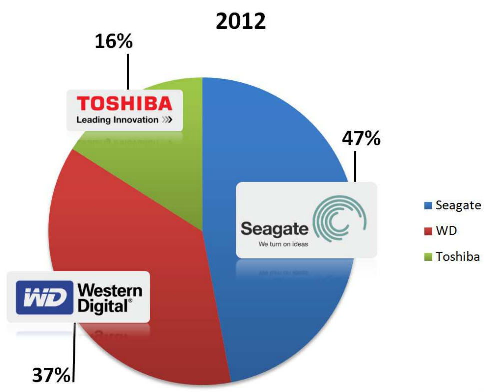


All drives today use the perpendicular magnetic recording technology (PMR) that was introduced already in 2006. The highest densities reached so far are
740 Gbit/in$^2$ for 2.5” hard disks and 620 Gbit/in$^2$ for 3.5” hard disks, which corresponds to about 1 Tbyte per platter

inside the disk. The expected limit for the PMR technology is around 1 Tbit/in$^2$.

The technology to reach much higher densities is called EAMR (Energy-Assisted-Magnetic-Recording) or specifically (Seagate) HAMR (Heat-Assisted Magnetic Recording). The theoretical limits are in this case 5-10 Tbit/in$^2$, which translates into 3.5" disk sizes of 30-60 Tbytes. Seagate has demonstrated lately the first HAMR prototypes [Link 46] in the lab reaching densities of about 1 Tbit/in$^2$. It is expected that the first major manufacturing of HDDs with HAMR will not start before 2016-2017, as the technology is very complicated and requires high level of upfront investments.

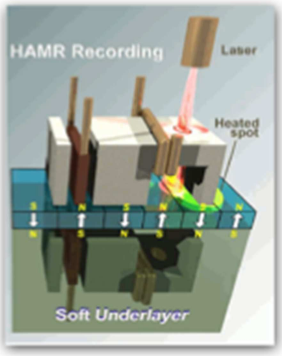


A more detailed roadmap of Western Digital can be found in this [Link 47] article.

The HDD industry is also considering an intermediate technology which would be able to reach ~2 Tbit/in$^2$ densities with relatively modest investment. But this Shingling Technology [Link 48] has some important consequences for the usage of HDDs. While the read performance is scaling with the density, the write performance drops considerably as all bits have to be written several times (3-4). In addition any random write performance will be severely degraded. Thus these type of disks could not be used for OS system disks but rather as write-once and write-in-one-go data disks. These disks would be used in a completely different market segment.

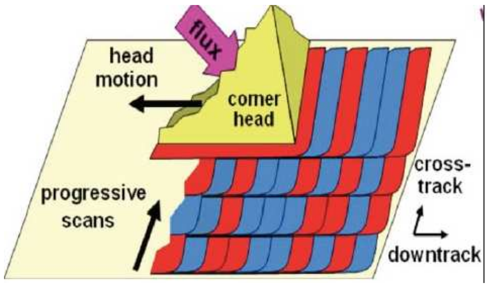


The price/performance evolution of consumer grade hard disks was following a clear exponential decrease over the last 20 years. Figure 15 shows this curve with data points added for 2012 [Link 49].

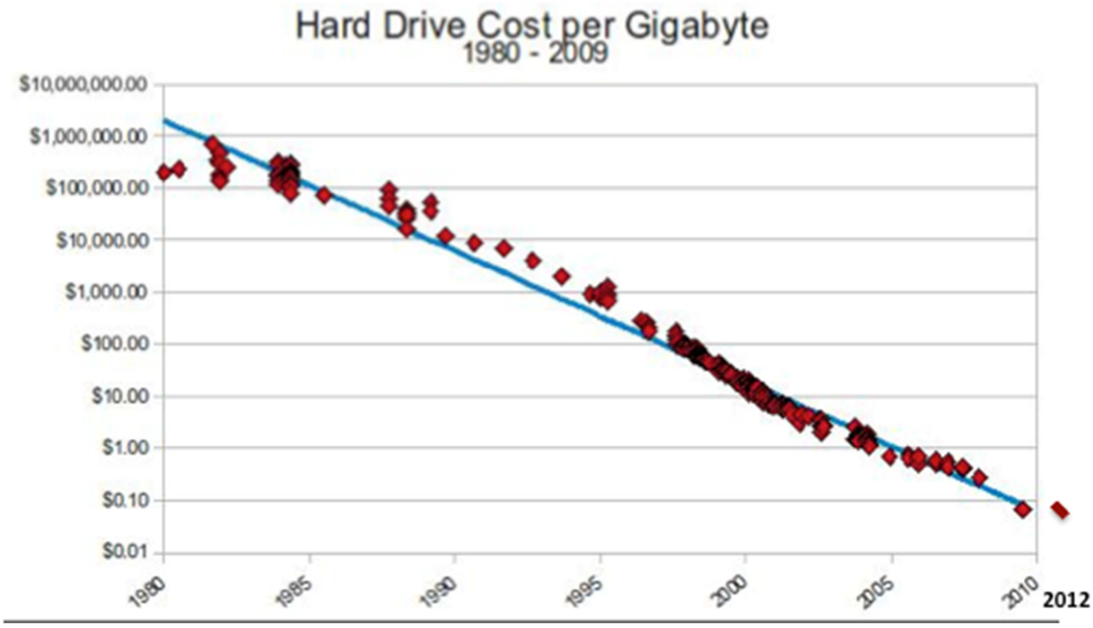


*Figure 15: Cost evolution of consumer grade HDD space*

This long term view shows a 60% price decrease per gigabyte per year.
But the picture changes considerably if one zooms into the details for the last few years. The curve has flattened and after the flood disaster in Thailand the prices have increased by a factor 2 and are only slowly approaching the values from mid 2011. [Link 50]
The flowing plot is based on street prices [Link 51] of enterprise type Hard-Disks (3 year warranty, large cache, 24*7 operation, etc.).

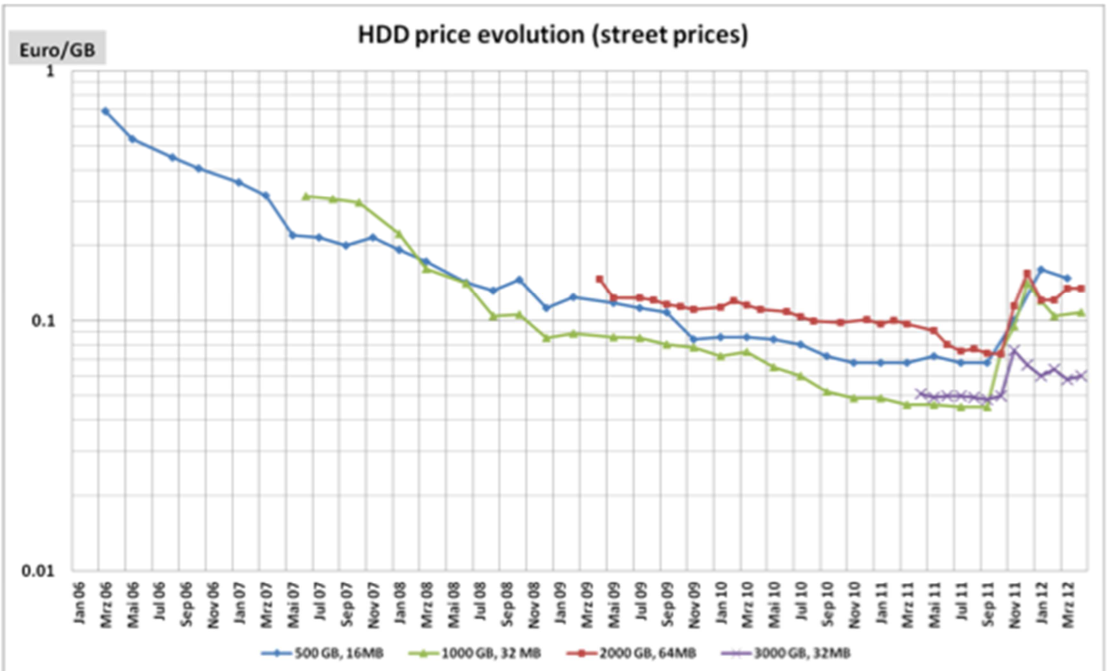


*Figure 16: Cost evolution of server quality HDD space*

There is a strong expectation that the price/performance improvement rate will slow down to about 20% per year. This is due to several factors:

- The current technology is reaching its limits and large investments are needed for the implementation of the next technology generation
- The demand for disk space is very high (50% from cloud storage) and the current production capacity is not able to fulfil this. The estimated demand in 2011 was 400 Exabyte, while 350 Exabytes were delivered.
- The price increase last year led to very high profits in the HDD industry

## 4.4. Tape storage

The total revenues in the tape market 2011 were about 3 billion $, shared more or less equally between tape drives, tape media and robotics. All three areas experienced a considerable negative growth rate (-10%), as the low and medium size backup market is moving to disk-based systems and the backup data size has been reduced by sophisticated methods (de-duplication, retention time optimization, etc.). The remaining market is concentrating on large-scale data storage and cloud storage. Figure 17 shows the expected need for archive storage during the next years.
[Link 52] [Link 53] [Link 54] [Link 55]

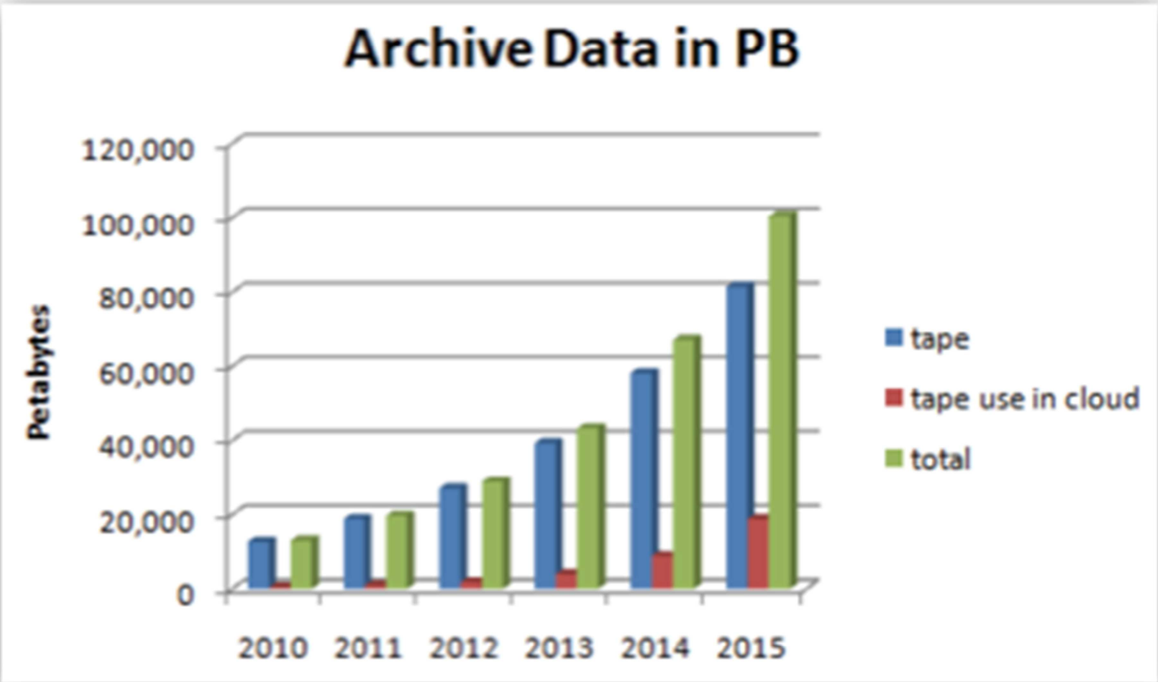


***Figure 17: Expected evolution of tape storage until 2015***

The LTO consortium is still dominating the market with a share of 90%.
The updated LTO roadmap extends now to LTO-8 in 2016. [Link 56]

| Feature / Model | LTO-3 | LTO-4 | LTO-5 | LTO-6 | LTO-7 | LTO-8 |
|---|---|---|---|---|---|---|
| Availability | Shipping | Shipping | Shipping | 2012 | 2014 | 2016 |
| Native Capacity | 400GB | 800GB | 1.5TB | 3.2TB | 6.4TB | 12.8TB |
| Capacity w/Compression | 800GB | 1.6TB | 3.0TB | 8TB | 16TB | 32TB |
| Native Data Rate (MB/s) | 80 | 120 | 140 | 210 | 315 | 472 |
| Data Rate w/Compression | 160 | 240 | 280 | 525 | 788 | 1,180 |
| WORM Support | Yes | Yes | Yes | Yes | Yes | Yes |
| Encryption Built-in | No | Yes | Yes | Yes | Yes | Yes |
| Interfaces Supported | LVD, FC | LVD, FC | FC, SAS | FC, SAS | FC, SAS | FC, SAS |

***Table 5: Characteristics of the various LTO tape generations***

Last year there have been some major tape technology upgrades in the enterprise tape area. IBM and Oracle increased their tape capacity to 4.0 TB and 5.0 TB, respectively (IBM 3592JC and Oracle T10000T2). Tape performance is now >200 MB/s. Both vendors have roadmaps for the next years.

The cost of tape media has fallen to a level of 0.04 CHF/GB. Figure X shows the cost evolution of LTO media over the last years and a prediction for the next 3 years. [Link 52]

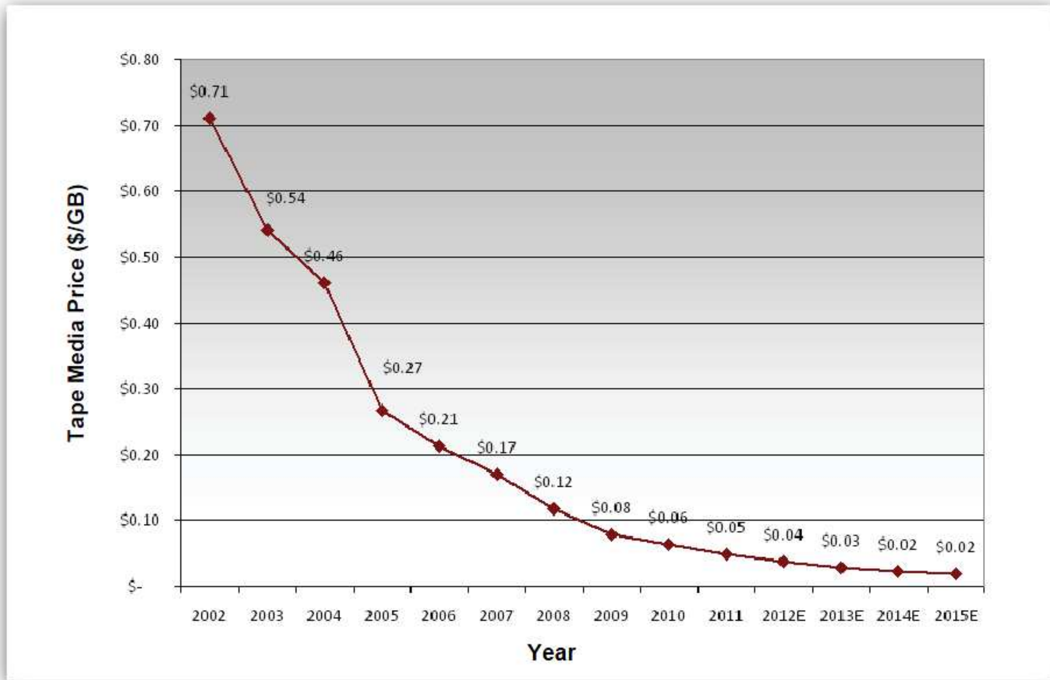


***Figure 18: Evolution of the tape space costs (media costs only)***

Especially with the price increase of disk space and the much better power footprint, tape storage still offers the best price/storage/power ratio for large scale data archiving.

# 5. Server performance and costs

The typical CERN CPU server configuration is the following:

- Dual socket server motherboard, half-size
- Two Intel Xeon server processors, low power series
- 2-3 GB of memory per physical core
- 2-3 local disks for CPU server (HDD 2 TB or SSD 250 GB)
- 2u chassis for 4 CPU sever motherboards, redundant power supply setup for 4 nodes and 1 Gbit Ethernet network connection

For storage servers we have now two different configuration types:

1. 4u chassis with dual-socket motherboard, single Intel Xeon server L-series processor, 12 GB memory, 24 or 36 3 TB SATA II enterprise HDDs, 10 Gbit ethernet connection
2. A CPU server node with one 10 Gbit connection and one SAS connector PCIe board plus one 36-bay external storage array with 36 3TB SATA II enterprise level HDDs

The details of these configurations are determined by several boundary conditions: Limited power and cooling in the computer centre, power-density space limits, requirements from the physics applications and cost-budget constraints.
The following measurements are based on the servers installed in the CERN computer centre over the last 7 years. Figure 19 shows the evolution of the overall performance of CPU servers. The performance is measured in HS06 (HepSpec 2006) based on a variant of the Spec benchmarks [Link 57].

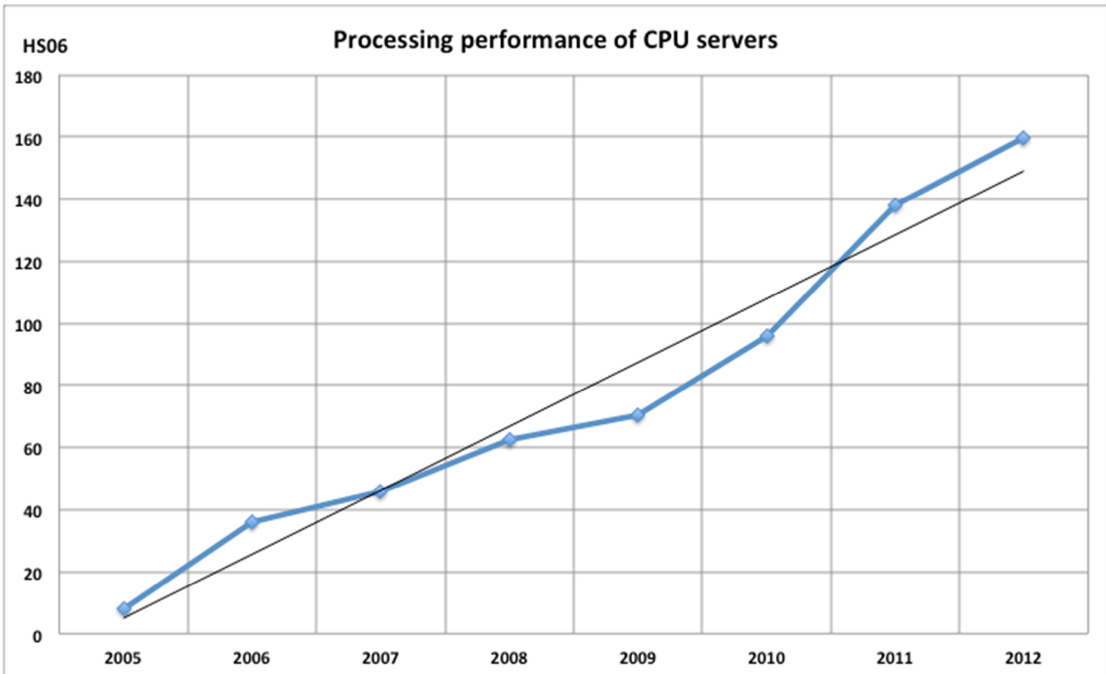


***Figure 19: Evolution of the processing performance of CERN CPU servers***

The increased performance per server and year follows clearly a linear curve. The same is true for the performance evolution of a single physical core as part of the used processors (Figure 20).

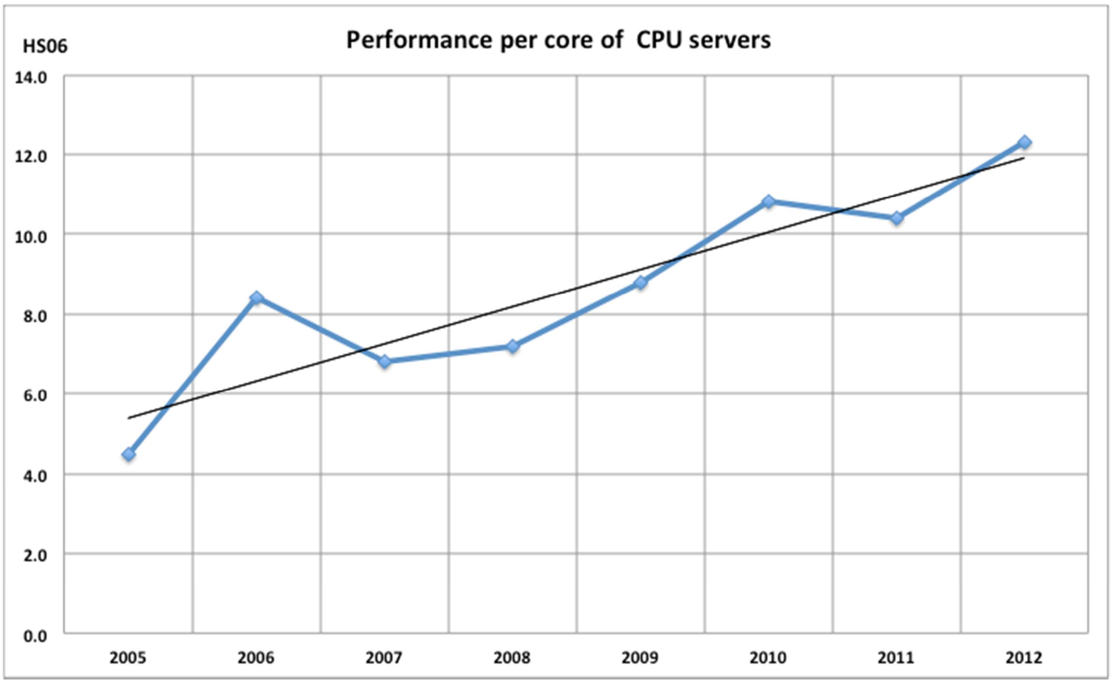


***Figure 20: Evolution of the processing performance of physical cores in CERN CPU servers***

The power efficiency of the CPU servers has steadily improved over the last years, but is again following a linear improvement curve.

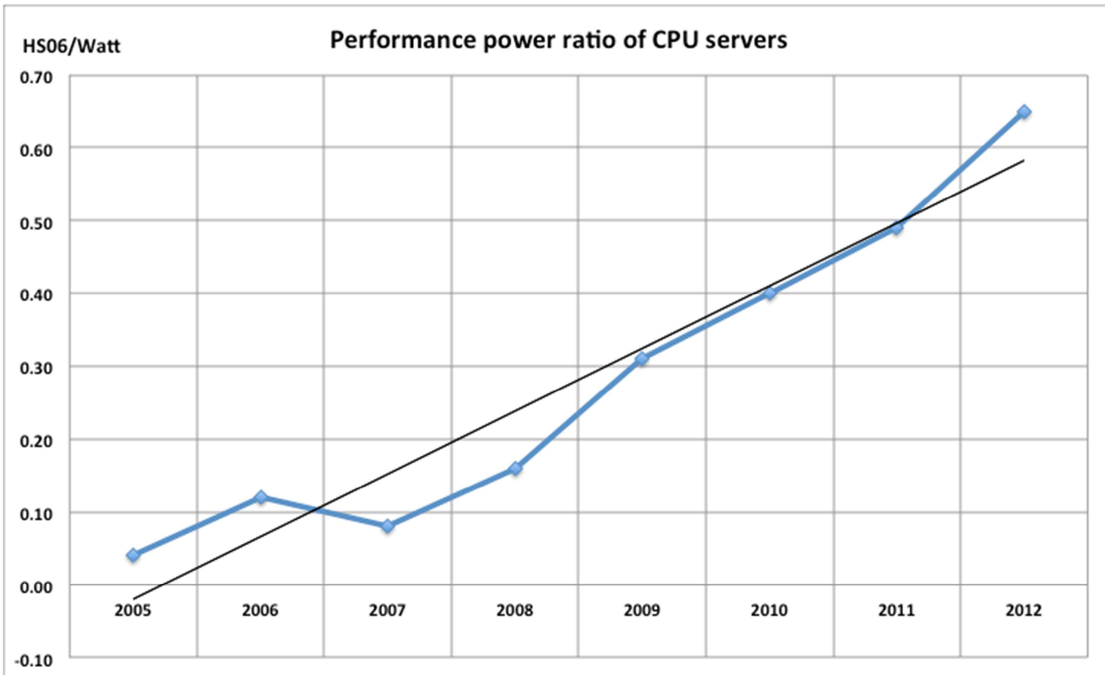


***Figure 21: Evolution of the CPU processing performance per electrical power consumption of CERN CPU server***

Similar improvement can be observed for the storage servers as shown in figure 22. This is of course closely related to the areal density improvements of HDDs, but weakened by the fact that disk severs need a certain constant infrastructure (motherboard, processors, memory) which follow different improvement curves.

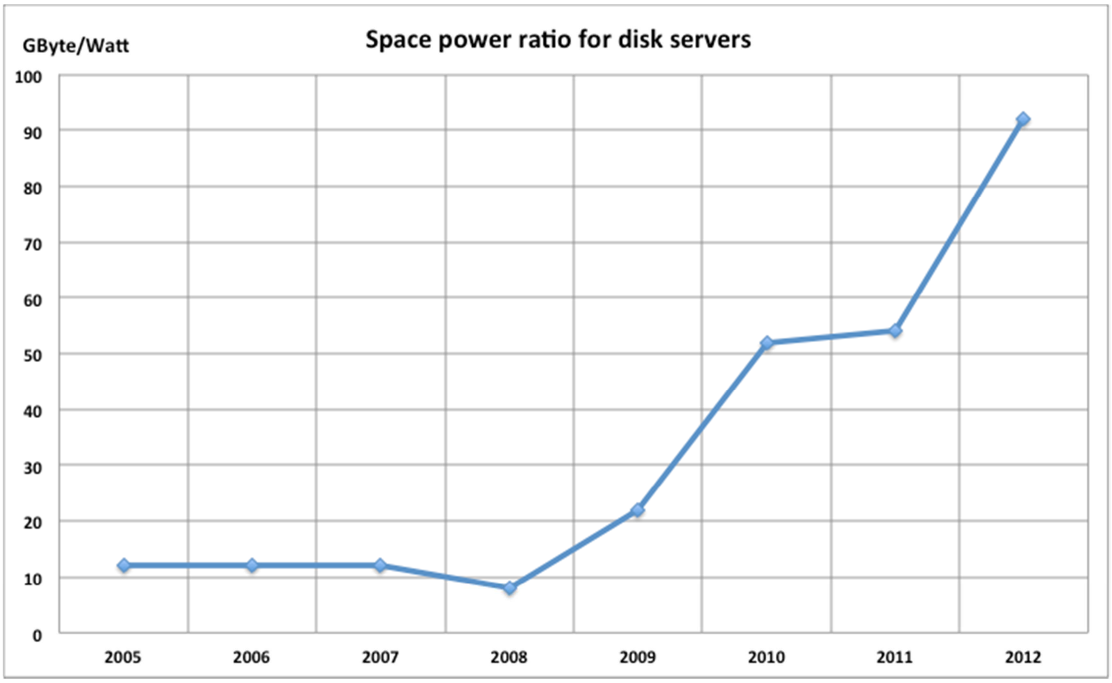


***Figure 22: Evolution of the storage capacity per electrical power consumption of CERN disk server***

The various quoted price/performance comparisons and evolutions (Figure , Figure, Figure ) were all done in Euro, while the cost predictions for the CPU and storage serves will be done in Swiss-Franc. To be able to convert the cost of computing equipment between the two currencies, one has to take into account two sets of information:

1. CERN is purchasing all equipment in Swiss-Franc and does not pay any VAT. The average VAT value in the main CERN member states was and is about 20%

2. The Euro Swiss-franc exchange rate has varied quite a bit during the last years as shown in Figure 23


***Figure 23: Evolution of the Euro-Swiss Franc exchange rate***

The cost calculations and predictions of CPU and disk servers depends of course on the price/performance evolution of the three main server components: processors, memory and disks. The cost of motherboards, chassis and power-supply can be considered as a relatively stable infrastructure base cost.
The street price development of memory (Figure 12), HDDS (Figure 16) and processors (Figure 24) shows clearly that these are not following simple mathematical curves. There is a general price drop with each new technology generation, some intermediate continuous price decrease over time and unexpected price increases.

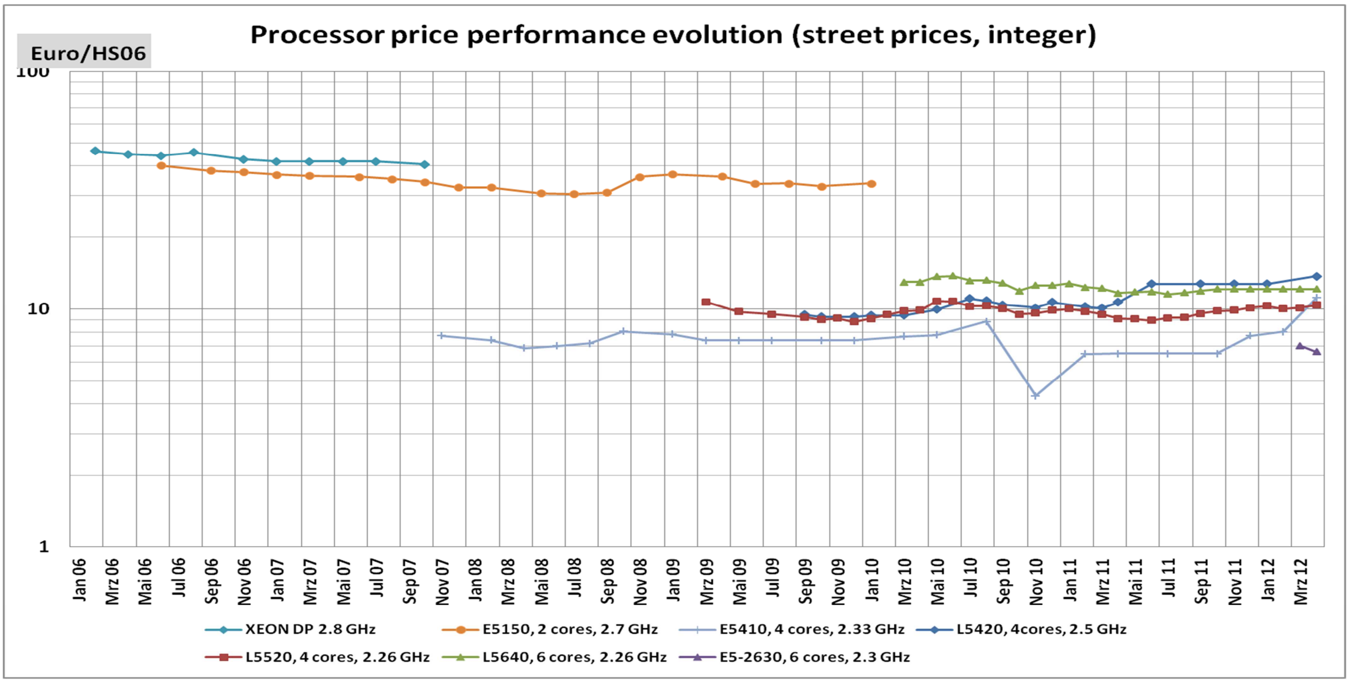


***Figure 24: Street price evolution of Intel Xeon server processors***
Street prices sources of memory, HDDs and processors [Link 51]

Figure 25 shows the price/performance development of CERN CPU server over the last few years and a scenario for the next 5 years. The values are normalized to CPU servers with 2 GB of memory per physical core.

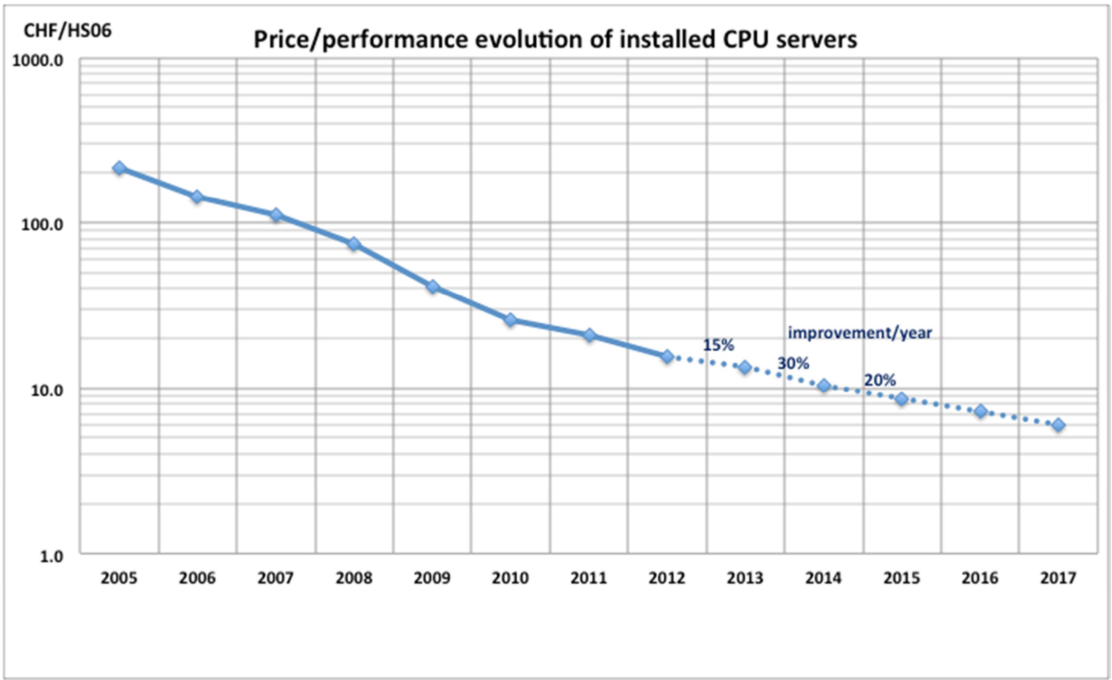


*Figure 25: Evolution and prediction of the price/performance of CPU servers*

The 2010 predictions for CPU server costs assumed a 40% improvement rate per year, but it is clear that this was over-optimistic as can be seen from the following table. One special point was the large increase of memory prices in 2010/2011.

| | 2011 | 2012 | 2013 | 2014 | 2015 | 2016 | 2017 |
|---|---|---|---|---|---|---|---|
| Jan 2010 prediction CHF/HS06 | 18.6 | 13.3 | 9.5 | 6.8 | 4.8 | | |
| May 2012 prediction CHF/HS06 | 21.0 | 15.5 | 13.5 | 10.4 | 8.6 | 7.2 | 6.0 |

*Table 6: Price/performance predictions (old and new) for CPU servers*

The new predictions are based on the extrapolation of the price/performance measurements from the last two years and incorporate the current market and technology trends:

- The prediction for 2013 is only a 15% price decrease, as the Sandy Bridge server processors have only been introduced right now and from figure 24 one can see that processor prices stay more or less constant within on technology line.

- One expected the next generation Ivy Bridge server processors to enter the market in mid 2013, thus the 30% improvement for 2014.
- It is expected that the move to 14nm structure sizes in >2015 will be technology-wise very challenging and also very expensive, thus the slowdown of improvement to 20% from 2015 onwards

Figure 26 shows the price/performance development of CERN disk server over the last few years and a scenario for the next 5 years. The values are normalized to servers with a RAID 0 configuration for all disks within a server (internal tray or externally SAS attached).

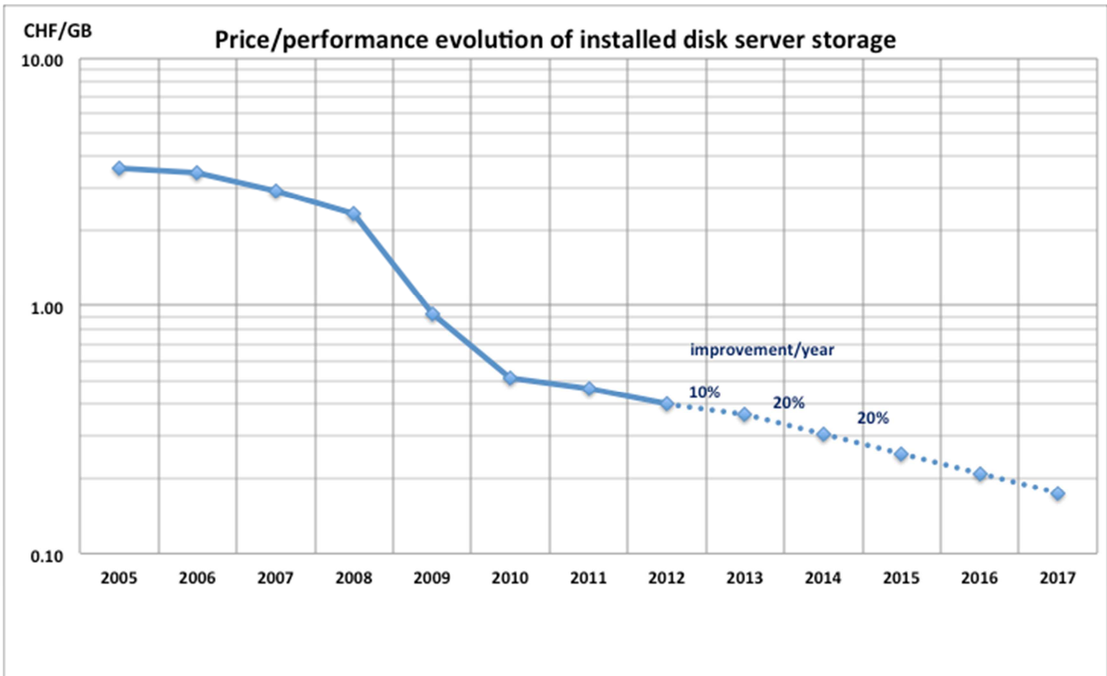


***Figure 26: Evolution and prediction of the price/performance of disk servers***

The 2010 predictions for disk server costs assumed a 35% improvement rate per year, but it is clear that this was over-optimistic as can be seen from the following table. The sharp cost decrease in 2008-2010 was due to the fast doubling of disk capacity from 0.5TB/disk to 2 TB/ disk in 3 years. It took more time to establish the 3 TB disk in the market and one is reaching the limits of the HDD PMR technology.

| | 2011 | 2012 | 2013 | 2014 | 2015 | 2016 | 2017 |
|---|---|---|---|---|---|---|---|
| Jan 2010 prediction CHF/GB | 0.37 | 0.27 | 0.20 | 0.15 | 0.11 | | |
| May 2012 prediction CHF/GB | 0.46 | 0.40 | 0.36 | 0.30 | 0.25 | 0.21 | 0.18 |

***Table 7: Price/performance predictions (old and new) for disk servers***

The new predictions are based on the extrapolation of the price/performance measurements from the last two years and incorporate the current market and technology trends:

- The HDD prices have not yet reached the level of mid 2011, thus the improvements for 2013 will only be 10%

- Form 2014 onwards the price improvements will be at a moderate 20% level, as the industry has to cope with the high demand of storage and the cost and problems of the new HAMR technology

## 6. New technologies

There is a staggering amount of research activity in the area of storage and processing. New material and methods to create transistors and store bits is reported nearly on a weekly basis. I have just listed a few of these research activities as an overview. I also incorporated results, which are already a few years old, but still very actively pursued. This gives an indication of how lengthy the time can be between basic research results and multi-million unit production. A good example for a complete failure is holographic storage.

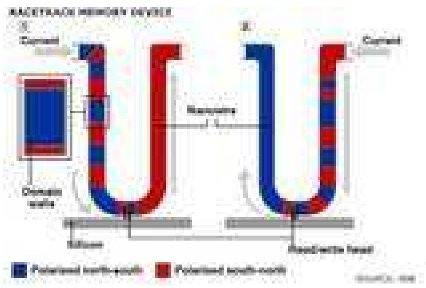


**IBM:** racetrack memory; moving domain walls in a nanowire; already from 2008 [Link 59] [Link 60]

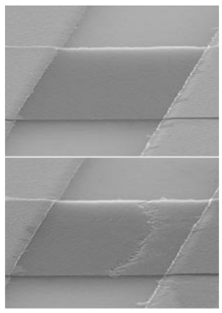

**Rice-University:** 3D graphene storage based on creating cracks; already from 2009 [Link 61]

These cracks can alos be re-used to create special structures [Link 62]

**Samsung:** Combine graphene and silicon to create transistors

**University Singapure:** Combining graphene with ferroelectrics for storage devices

**University California/Taiwan:** Embedded 3nm silicon nanodots [Link 63]

**IBM:** Atomic scale magnetic computer memory [Link 64]

**Memristor:** First production quality products expected in mid-2013
2008 predictions: replacing Flash in 2012, DRAM in 2014, HDD in 2016 !
[Link 60]

**PCRAM** Phase change memory; 8Gbit 20nm PCR prototype device was produced by Samsung in Nov 2011 [Link 65]

**MRAM** already in production state,; 16Mbit chips, 130nm structure size, non-volatile, 50ns latency, ~3 million units shipped in 2011 [Link 66]
The 2005 predictions were assuming a multi-billion market , real market is < 50 million

Single phosphorus atom transistor on silicon
[Link 67]

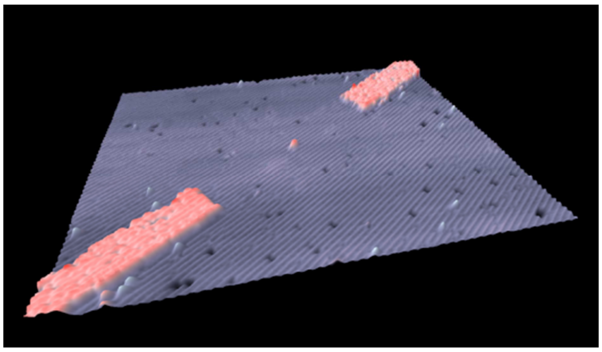

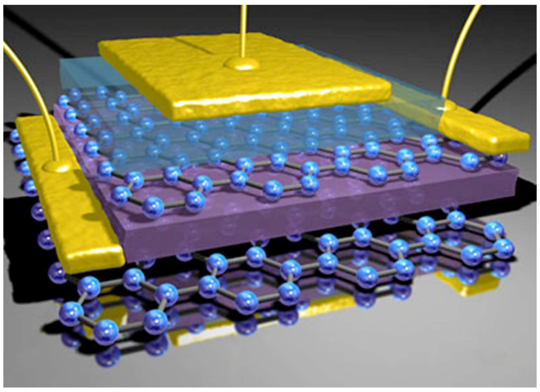

Single atom layered transistors based on grapheme and boron nitride [Link 68]

9nm transistor based on nanotubes [Link 69]

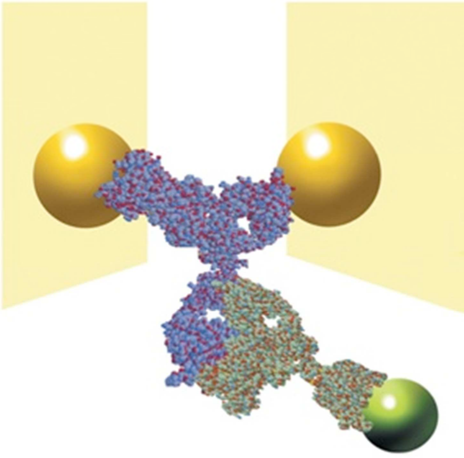

Self-assembly of transistors based on biological molecules (Immunoglobulin G antibody) and 5nm gold particles
[Link 70]

The atomic-thin layers of the carbon allotrope graphene are a key research focus for new was to implement electronic components. Lately it has been discovered that a chemical component based on the metal molybdenum shows very similar electronic characteristics and could become a 'rival' of graphene. [Link 71]

But all these research activities have produced only small size test samples or prototypes in the laboratory. Moving these discoveries into commodity products, which are manufactured, on the multi-million or even billion scale is a very time consuming and expensive process. Figure 27 is an extract from the 2011 ITRS report (International Technology Roadmap for Semiconductors) [Link 72] describing their estimate of when several new semiconductor technologies could be used in larger scale manufacturing processes. A more detailed description of the estimates and problems can be found in an article on the ExtremTech website [Link 73].

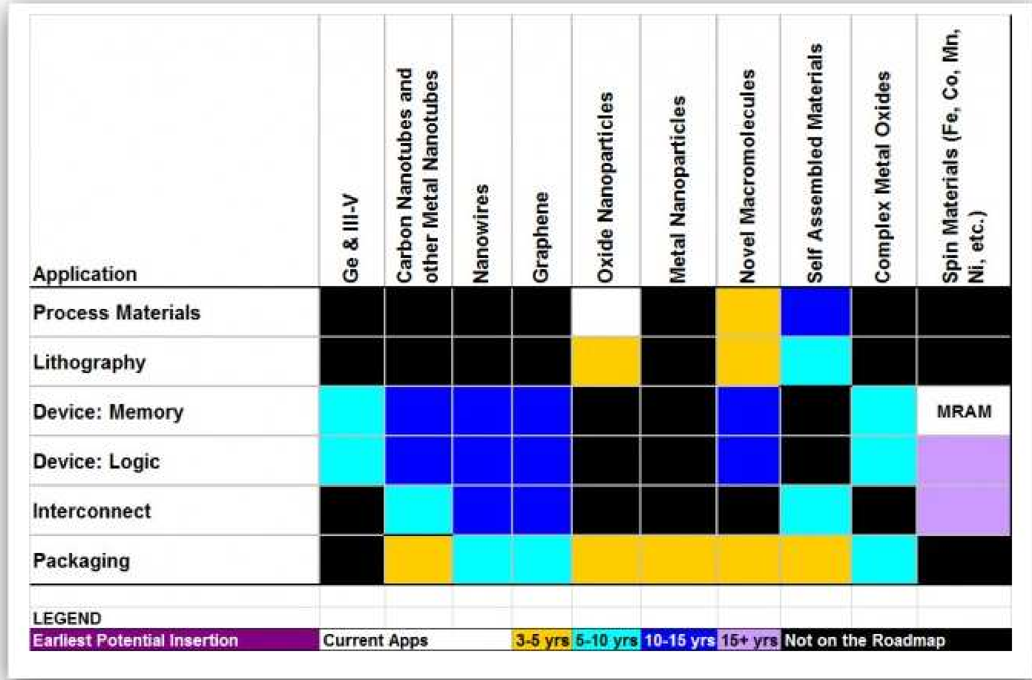


*Figure 27: Estimates for new semiconductor technologies to be moved into production*

# 7. Non-mainstream computing

There is quite a range of unconventional computing concepts and I am listing here only a small number of examples. Some are just prof-of-concepts others might have wide ranging consequences and a few are just amusing.

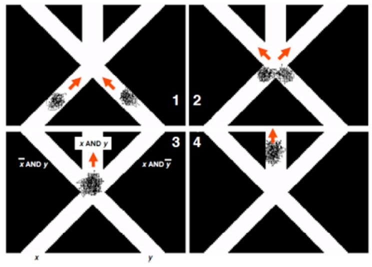


- Implementation of AND/NOR computing units based on the movement of hungry soldier crabs [Link 74]

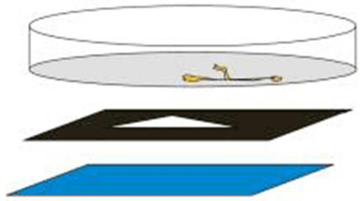

- Biocomputing using a light sensitive single-cell organism (plasmodium) e.g. find the shortest way through a maze [Link 75]

- Chemical computing with traveling wave fronts of crystallization in supersaturated sodium acetate (hot-ice) [Link 76]
- Molecular Cryptosystem for Images by DNA Computing [Link 77]
- DNA as a storage device [Link 78]
- General computing concepts implemented with DNA [Link 79]
- New processor concept doing 'inexact' computing to save energy [Link 80]

There are of course the regular 'breakthrough' research results in quantum computing: reaching 14 qbit systems; >100 us coherence time; stable qbits with seconds lifetime; moving from atomic gases to nano-dots for the quantum states, etc. But we are still at least 10 years away from any usable production state quantum computer.

Last but not least one has to mention the 'Internet of Things' which is somewhat coupled to the general proliferation of mobile devices and cloud computing plus storage. The technology of merging digital, analogue and sensor components into small and cost-effective chips is driving the 'Internet-of-things' initiatives, creating very large amounts of nano-data streams from ubiquitous sensors [Link 81].

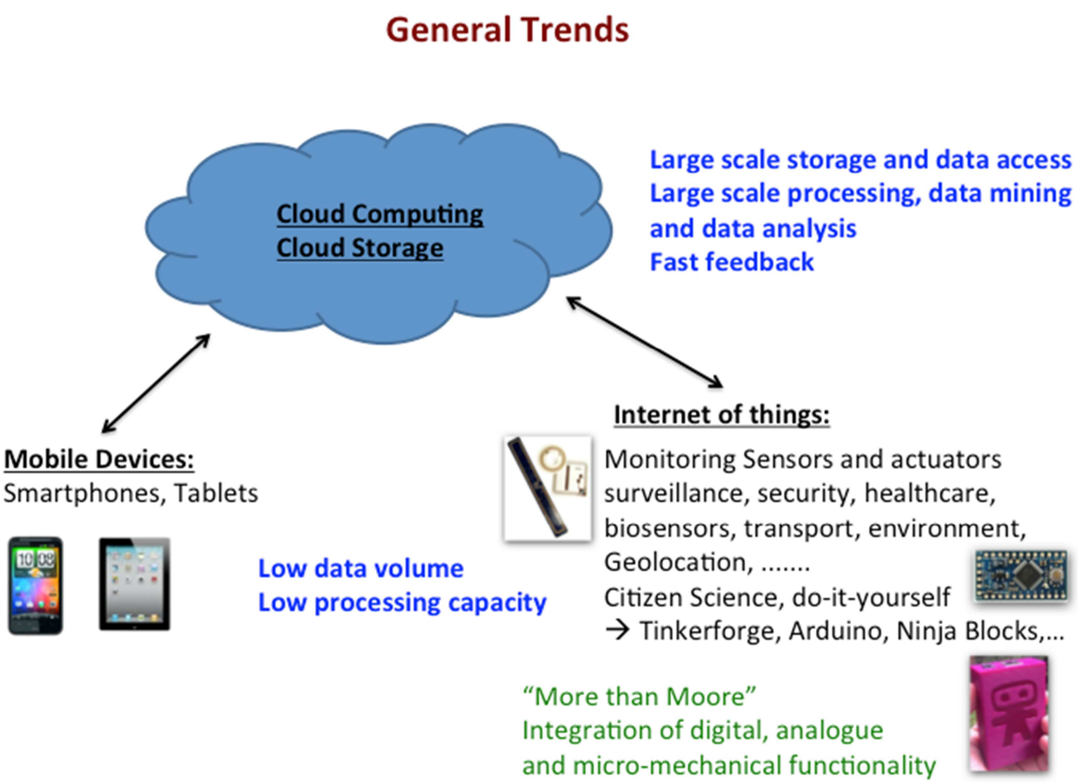


***Figure 28: Cloud computing, mobile devices and the Internet of Things***

# 8. Summary

The market and technology analysis presented in the previous chapters can be summarizes in a few key statements:

- There is a considerable move of the market towards energy efficient mobile devices. This goes along with a strong increase of cloud computing and storage services.
- Even in the server area energy efficiency has a stronger focus than performance
- There is a dominance of only a few companies in nearly all computing component markets (disk, CPU, NAND, tape, graphics).
- Chip complexity evolution is still following Moore's Law
- Other computing characteristics have started to follow linear instead of exponential growth rates (e.g. core count, energy efficiency)
- Moore's Law is about chip complexity and has a complicated relationship to actual calculations of server costs. The individual components of servers show very different price/performance behaviours with different slopes and sometimes large fluctuations.
- Technology roadmaps are well advanced, but approaching 10nm structure sizes will be challenging and very expensive (~2016).
- There is a constant flow of very exciting bleeding edge research results, but they are far away from having any market relevance.
- There is still a continuous positive evolution of the price/performance improvements of all relevant computing parts, but there are strong indications that the large improvements from previous years will not be achieved anymore in the future (slowdown of improvements).

# References

[Link 1] Semiconductor market , Intel share, revenues, list of companies
http://www.isuppli.com/Semiconductor-Value-Chain/News/Pages/Intels-Semiconductor-Market-Share-Surges-to-More-Than-10-Year-High-in-2011.aspx

[Link 2] IDC, smart device shipment predictions
http://www.idc.com/getdoc.jsp?containerId=prUS23398412

[Link 3] Dark Silicon
http://www.google.ch/url?sa=t&rct=j&q=&esrc=s&source=web&cd=1&ved=0CCgQFjAA&url=http%3A%2F%2Fwww.cs.utexas.edu%2Fftp%2Fdburger%2Fpapers%2FISCA11.pdf&ei=9D6ET7aoEcSmhAfU1cncCA&usg=AFQjCNHV0um4_Oxwunyqp6h6vfQvTERzPA

[Link 4] AMD new server strategy, no core increase, focus on APU, 5-7% server market share
http://www.hpcwire.com/hpcwire/2012-02-07/amd_unveils_new_strategy_for_server_silicon.html

[Link 5] ISSCC, International Solid-State Circuits Conference
http://isscc.org/trends/index.html

[Link 6] Wafer sizes, industry roadmaps
http://chipsecurity.org/2012/03/450mm-wafer-size-change-and-its-security-implications/

[Link 7] Chip manufacturing, 90-22 nm, technology , tri-gate
http://semiaccurate.com/2012/05/30/intel-22nm-process-problems-part-2/
[Link 8] Intel server architecture roadmap
http://news.softpedia.com/news/Intel-2012-2018-Server-CPU-Roadmap-Reveled-213265.shtml

[Link 9] Intel details 10nm, 7nm, 5nm process roadmap
http://www.bit-tech.net/news/hardware/2012/05/14/intel-process-roadmap/1
http://semiconductorexpert.blogspot.ch/2012/05/intel-roadmap-to-2015-and-beyond-5nm.html

[Link 10] Extreme ultraviolet Lithography, description and problems
http://spectrum.ieee.org/semiconductors/devices/euv-faces-its-most-critical-test

[Link 11] ExtremeTech article, death of the cpu scaling
http://www.extremetech.com/computing/116561-the-death-of-cpu-scaling-from-one-core-to-many-and-why-were-still-stuck

[Link 12] Nvidia Kepler graphics processor
http://www.theregister.co.uk/2012/05/15/nvidia_kepler_tesla_gpu_revealed/

[Link 13] Cloud gaming, Onlive, Gaikai , Nvidia
http://www.onlivespot.net/2012/05/nvidia-unveiled-geforce-grid-cloud.html

[Link 14] Tilera, many-core system
http://www.hpcwire.com/hpcwire/2012-02-01/tilera_chips_in_with_64-bit_manycore_processors.html

[Link 15] Adapteva, many-core system
http://www.rethink-wireless.com/2012/03/20/adapteva-pushes-64-core-processor-smartphones.htm

[Link 16] ZiiLabs, many-core system
http://www.engadget.com/2012/01/05/ziilabs-debuts-100-core-zms-40-processor-optimized-for-android/

[Link 17] Fudan University, many-core system
http://semiaccurate.com/2012/02/21/chinese-16-core-cpu-uses-message-passing/

[Link 18] Plurality, man-core system
http://www.plurality.com/hypercore.html

[Link 19] Neuromorphic processors
http://electronics.ezinemark.com/neuromorphic-third-eye-sensors-7d34ef85dbe2.html

[Link 20] Neuromorphic processors
http://celest.bu.edu/outreach-and-impacts/the-synapse-project

[Link 21] Kalray, many-core processor
http://www.eetimes.com/electronics-news/4235118/European-startup-attempts-many-core-revolution

[Link 22] Intel, MIC co-processor
http://www.drdobbs.com/parallel/232800139

[Link 23] Enterpoint, FPGA co-processor boards
http://www.hpcwire.com/hpcwire/2012-03-24/enterpoint_offers_fpga_card_for_desktop_mobile_hpc.html

[Link 24] Maxeler supercomputer, JP Morgan
http://www.forbes.com/sites/tomgroenfeldt/2012/03/20/supercomputer-manages-fixed-income-risk-at-jpmorgan/

[Link 25] BittWare, FPGA co-processor boards
http://www.bittware.com/media/press/pr.cfm?id=70

[Link 26] HPC tests with FPGAs
http://www.hpcwire.com/hpcwire/2012-04-16/latest_fpgas_show_big_gains_in_floating_point_performance.html

[Link 27] Calxeda, micro server with ARM processors
http://www.calxeda.com/company/newsroom/pressreleases/63-calxeda-launches-energycore-processor

[Link 28] HP project Moonshot, micro server with ARM processors
http://www.slashgear.com/hp-project-moonshot-reveals-low-power-redstone-arm-servers-02192452/

[Link 29] ARM processors for HPC
http://www.hpcwire.com/hpcwire/2011-11-09/arm_yourselves_for_exascale,_part_1.html
http://www.hpcwire.com/hpcwire/2011-12-12/arm_yourselves_for_exascale,_part_2.html

[Link 30] SeaMicro, micro server
http://www.theregister.co.uk/2012/03/05/seamicro_amd/

[Link 31] Article on multi and many cores problems
http://www.extremetech.com/computing/116561-the-death-of-cpu-scaling-from-one-core-to-many-and-why-were-still-stuck/3

[Link 32] Dongarra, top500, software is the problem, hardware limits
http://www.netlib.org/utk/people/JackDongarra/SLIDES/uu-sci-0212.pdf

[Link 33] IBM Renaissance Project, programming languages
http://my-inner-voice.blogspot.com/2012/01/many-core-processors-everything-you.html

[Link 34] IBM RFontana, April 2012. Technology roadmap comparison, tape, HDD, NAND
http://storageconference.org/2012/Presentations/M09.Fontana.pdf

[Link 35] DRAM roadmap, JEDEC
http://www.jedec.org/category/technology-focus-area/main-memory-ddr3-ddr4-sdram

[Link 36] DRAM technology roadmap
http://semimd.com/blog/2011/11/03/dram-server-roadmap-points-to-ddr4-and-3d/

[Link 37] Memory wall
http://denalimemoryreport.wordpress.com/2012/04/16/will-your-multicore-soc-hit-the-memory-wall-will-the-memory-wall-hit-your-soc-does-it-matter/

[Link 38] DRAM 3D chip stacking
http://www.extremetech.com/computing/119843-the-future-of-computers-3d-chip-stacking

[Link 39] Hyper Memory Cube
http://www.electroiq.com/blogs/insights_from_leading_edge/2011/11/iftle-74-the-micron-memory-cube-consortium.html

[Link 40] Global NAND market figures
http://www.4-traders.com/IHS-INC-13057/news/IHS-NAND-Flash-Market-To-Grow-8-This-Year-On-Ultrabook-Demand-14282053/

[Link 41] SanDisk announcement, 19nm flash chip, 128gb, three bit MLC
http://www.engadget.com/2012/02/22/sandisk-makes-128-gigabit-flash-chip-crams-three-bits-per-cell/

[Link 42] Endurance problems of NAND memory
http://www.anandtech.com/show/5067/understanding-tlc-nand/
[Link 43] Bleak future of NAND memory

http://www.tomshardware.com/news/NAND-Flash-SSD-MLC-TLC-Laura-Grupp,14728.html

[Link 44] Global SSD market figures
http://www.digikey.com/purchasingpro/us/en/articles/semiconductors/ultrabooks-ipads-and-iphones-will-drive-nand-flash-growth-in-2012/1328

[Link 45] HDD market consolidation
http://news.techeye.net/business/hdd-business-to-become-mexican-standoff

[Link 46] Seagate announcement, HAMR 1 TB/in2
http://www.theregister.co.uk/2012/03/20/seagate_terabit_areal_density/

[Link 47] HDD Western Digital roadmaps
http://www.theregister.co.uk/2012/05/09/wd_disk_tech_views/

[Link 48] HDD intermediate technology, shingling
http://www.theregister.co.uk/2011/07/29/hgst_smr/

[Link 49] HDD cost evolution for consumer grade disks
http://ns1758.ca/winch/winchest.html

[Link 50] HDD price evolution 2012
http://www.extremetech.com/computing/129874-hdd-pricewatch-higher-prices-are-the-new-normal

[Link 51] Computing equipment street price references
http://skinflint.co.uk/?m=1

[Link 52] International magnetic tape storage roadmap INSIC
http://www.insic.org/news/2012Roadmap/news_12index.html

[Link 53] Tape roadmaps
http://www.storagenewsletter.com/news/tapes/insic-tape-roadmap
[Link 54] Cloud providers and tape
http://www.theregister.co.uk/2011/11/02/spectralogic_t_finity/

[Link 55] Blue waters supercomputer and tapes
http://www.theregister.co.uk/2012/05/22/blue_waters_spectra_logic/

[Link 56] 2012 LTO roadmap
http://www.qualstar.com/146252.html

[Link 57] HEP spec benchmark
https://twiki.cern.ch/twiki/bin/view/FIOgroup/TsiBenchHEPSPEC

[Link 58] Euro to Swiss Franc exchange rate
http://www.finanztreff.de/kurse_einzelkurs_charts.htn?i=2203841&zeit=10000 0&sektion=zoom&popup=popup

[Link 59] IBM racetrack memory
http://www.dailytech.com/IBMs+New+Racetrack+Memory+is+on+an+IC+But+ Only+has+One+Car+Per+Lane/article23442.htm

[Link 60] IBM racetrack memory and memristor
http://www.techpowerup.com/156317/R-amp-D-IBM-s-Racetrack-Memory-Data-Storage-At-Superfast-DRAM-Speeds.html

[Link 61] 3D graphene storage
http://www.technologyreview.com/computing/23468/

[Link 62] Graphene cracks creating structures
http://biygroup.blogs.rice.edu/archives/155

[Link 63] Silicon nanodots
http://www.theregister.co.uk/2012/04/20/silicon_nanodot_memory/

[Link 64] IBM atomic scale memory
http://www.technologyreview.com/computing/39450/

[Link 65] Samsung PCRAM prototype
http://www.eetimes.com/electronics-news/4230958/ISSCC--Samsung-preps-8-Gbit-phase-change-memory

[Link 66] MRAM production and state
http://www.mram-info.com/tags/mram_production

[Link 67] Single phosphorus atom transistor
http://www.nature.com/nnano/journal/v7/n4/full/nnano.2012.21.html

[Link 68] Single atom layered transistors
http://physicsworld.com/cws/article/news/2012/apr/02/electrons-tunnel-nicely-through-boron-nitride

[Link 69] Nanotube transistor http://www.bit-tech.net/news/2012/01/30/ibm-carbon-nanotube/1

[Link 70] Self-assembly transistors
http://www.nature.com/nnano/journal/v7/n3/full/nnano.2012.7.html

[Link 71] Molybdenum disulifde as a possible graphene rival
http://spectrum.ieee.org/semiconductors/nanotechnology/graphenes-new-rival

[Link 72] ITRS report, International Technology Roadmap for Semiconductors, 2011
http://www.itrs.net/Links/2011ITRS/2011Chapters/2011ERM.pdf

[Link 73] ExtremeTech article, cutting edge technologies, cpu, citing from the ITRS report
http://www.extremetech.com/extreme/120353-the-future-of-cpu-scaling-exploring-options-on-the-cutting-edge/2

[Link 74] Computing with crabs
http://www.technologyreview.com/blog/arxiv/27730/?p1=blogs

[Link 75] Computing with plasmodium
http://www.technologyreview.com/blog/arxiv/23962/

[Link 76] Chemical computing example
http://www.technologyreview.com/blog/arxiv/24072/

[Link 77] DNA computing example, encryption
http://onlinelibrary.wiley.com/doi/10.1002/anie.201107156/abstract

[Link 78] DNA computing example, storage
http://www.pnas.org/content/early/2012/05/14/1202344109.abstract?sid=a8c9987b-03b5-49ac-834c-dd213e74f74a

[Link 79] DNA computing overview
http://www.economist.com/node/21548488

[Link 80] inexact processors
http://www.futurity.org/top-stories/faulty-computer-chip-is-15x-more-efficient/

[Link 81] Ubiquitous sensors, streams of nano-data
http://www.technologyreview.com/business/40016/?p1=BI