

Report on Disruptive Technologies for years 2020-2030

Version August 2016

Authors:

Prof. Dr. Theo Ungerer, University of Augsburg
Prof. Dr.-Ing. Dietmar Fey, University of Erlangen-Nuremberg

Compiled by:

Mike Knebel, University of Augsburg

With contribution of:

Nader Bagherzadeh, University of California, Irvine – *3D stacking*

Sandro Bartolini, University of Siena - *photonics*

Koen Bertels, Delft University of Technology - *quantum computing*

Christian Bradatsch, University of Augsburg - *graphene*

Jose Manuel García Carrasco, University of Murcia - *photonics*

Koen De Bosschere, Ghent University - *overall comments*

Marc Duranton, CEA LIST DACLE - *various technologies*

Babak Falsafi, Ecole Polytechnique Federale de Lausanne -
various technologies

Martin Frieb, University of Augsburg – *CMOS scaling*

Florian Haas, University of Augsburg - *photonics*

Said Hamdioui, Delft University of Technology - *quantum and
resistive computing*

Florian Kluge, University of Augsburg - *3D stacking*

Mike Knebel, University of Augsburg - *diamond computing, overall
compilation*

Avi Mendelson, Technion - *diamond computing*

Jörg Mische, University of Augsburg - *memristors, resistive
computing*

Nizar Msadek, University of Augsburg - *quantum computing*

Benjamin Pfundt, University of Erlangen-Nuremberg - *3D
stacking, memristors, resistive computing*

Ulrich Rückert, University of Bielefeld - *neuromorphic computing*

Alexander Stegmeier, University of Augsburg - *nanotubes*

Sebastian Weis, University of Augsburg - *neuromorphic computing*

Abstract

This report is part of the roadmapping effort within the EC CSA Eurolab-4-HPC. The roadmap itself targets a long-term roadmap (2022-2030) for High-Performance Computing (HPC) and it was decided, because of the speculative nature, to start with an assessment of future computing technologies that could influence HPC hardware and software.

The report covers the following technologies: CMOS scaling, die stacking and 3D chip technologies, Non-volatile Memory (NVM) technologies, Photonics, Resistive Computing, Neuromorphic Computing, Quantum Computing, Nanotubes, Graphene and Diamond Transistors.

From the assessment of these technologies we derive some potential long-term impacts of Disruptive Technologies for HPC hardware.

The report is the draft (August 2016) of an on-going assessment process and will be extended in the future.

Table of contents of Appendix Report on Disruptive Technologies

1. INTRODUCTION	4
2. IMPACT OF DISRUPTIVE TECHNOLOGIES	5
<i>Summary of Potential Long-Term Impacts of Disruptive Technologies for HPC Hardware</i>	5
<i>Summary of Potential Long-Term Impacts of Disruptive Technologies for HPC Software and Applications</i>	6
3. SUSTAINING TECHNOLOGY (IMPROVING HPC HW IN WAYS THAT ARE GENERALLY EXPECTED)	9
<i>Continuous CMOS scaling</i>	9
References	9
<i>Die Stacking and 3D-Chip</i>	11
References	13
4. DISRUPTIVE TECHNOLOGY IN HARDWARE/VLSI (INNOVATION THAT CREATES A NEW LINE OF HPC HARDWARE SUPERSEDING EXISTING HPC TECHNIQUES)	15
<i>Non-volatile Memory (NVM) Technologies</i>	15
References	17
<i>Photonics</i>	19
References	20
5. DISRUPTIVE TECHNOLOGY (ALTERNATIVE WAYS OF COMPUTING)	21
<i>Resistive Computing</i>	21
References	22
<i>Neuromorphic Computing</i>	25
References	28
<i>Quantum Computing</i>	29
References	31
6. BEYOND CMOS	32
<i>Nanotubes</i>	32
References	32
<i>Graphene</i>	33
References	33
<i>Diamond Transistors</i>	35
References	35

1. Introduction

Roadmapping beyond the upcoming Exascale machines (2022-2030) is extremely speculative. The basic idea of Eurolab-4-HPC roadmap is therefore to assess potentially disruptive technologies and summarize its impacts on HPC hardware as IF .. THEN .. statements, i.e. IF disruptive technology will be available THEN potential impact on hardware could be.

To sort the different technologies we define Types of Innovation adapted to HPC as:

Sustaining: An innovation that does not principally affect existing HPC. An innovation that improves HPC hardware in ways that were generally expected.

Discontinuous: An innovation that is unexpected, but nevertheless does not affect existing HPC.

Disruptive: An innovation that creates a new line of HPC hardware by applying a different set of values, which ultimately (and unexpectedly) overtakes existing HPC techniques.

We survey the current state of research and development and its potential for the future of the following hardware technologies:

- CMOS scaling
- Die stacking and 3D chip technologies
- Non-volatile Memory (NVM) technologies
- Photonics
- Resistive Computing
- Neuromorphic Computing
- Quantum Computing
- Nanotubes
- Graphene and
- Diamond Transistors

We categorize the technologies as:

- Sustaining technologies: CMOS scaling and Die stacking, see section 3
- Disruptive technologies that potentially create a new line of HPC hardware: NVM and Photonics, see section 4
- Disruptive technologies that potentially create alternative ways of computing: Resistive, Neuromorphic, and Quantum Computing, see section 5
- Disruptive technologies that potentially replace CMOS for processor logic: Nanotube, Graphene, and Diamond technologies, see section 6.

We summarize potential long-term impacts of Disruptive Technologies on HPC hardware in section 2 of the preliminary roadmap. Such impacts could concern the processor logic, the memory hierarchy, and potential hardware accelerators.

2. Impact of Disruptive Technologies

Summary of Potential Long-Term Impacts of Disruptive Technologies for HPC Hardware

Potential long-term impacts of disruptive technologies could concern the processor logic, memory hierarchy, and future hardware accelerators.

Processor logic could be totally different if materials like graphene, nanotube or diamond would replace classical integrated circuits based on silicon transistors, or could integrate effectively with traditional CMOS technology to overcome its current major limitations like limited clock rates and heat dissipation.

A physical property that these materials share is the high thermal conductivity: Diamonds for instance can be used as a replacement for silicon, allowing diamond based transistors with excellent electrical characteristics. Graphene and nanotubes are highly electrically conductive and could allow a reduced amount of heat generated because of the lower dissipation power, which makes them more energy efficient. With the help of those good properties, less heat in the critical spots would be expected which allows much higher clock rates and highly integrated packages. Whether such new technologies will be suitable for computing in the next decade is very speculative.

Furthermore, Photonics, a technology that uses photons for communication, can be used to replace communication busses to enable a new form of inter- and intra-chip communication.

Current CMOS technology may presumably scale continuously in the next decade, down to 6 or 5 nm. However, scaling CMOS technology leads to steadily increasing costs per transistor, power consumption, and to less reliability. Die stacking could result in 3D many-core microprocessors with reduced intra core wire length, enabling high transfer bandwidths, lower latencies and reduced communication power consumption.

3D stacking will also be used to scale flash memories, because 2D NAND flash technology does not further scale. In the long run even 3D flash memories will probably be replaced by memristor or other non-volatile memory (NVM) technologies. These, depending on the actual type, allow higher structural density, less leakage power, faster read- and write access, more endurance and can nevertheless be more cost efficient.

However, the whole memory hierarchy may change in the upcoming decade. DRAM scaling will only continue with new technologies, in fact NVMs, which will deliver non-volatile memory potentially replacing or being used in addition to DRAM. Some new non-volatile memory technologies could even be integrated on-chip with the microprocessor cores and offer orders of magnitude faster read/write accesses and also much higher endurances than flash. Intel demonstrated the possible fast memory accesses of the 3D-XPoint NVM Technology used in their Optane Technology. HP's computer architecture proposal called "The Machine" targets a machine based on new NVM memory and photonic busses. The Machine sees the memory instead of processors in the centre. This so called Memory-Driven Computing unifies the memory

and storage into one vast pool of memory. HP proposes advanced photonic fabric to connect the memory and processors. Using light instead of electricity is the key to rapidly accessing any part of the massive memory pool while using much less energy.

The Machine is a first example of the new Storage-class Memory (SCM), i.e., a non-volatile memory technology in between memory and storage, which may enable new data access modes and protocols that are neither 'memory' nor 'storage'. It would particularly increase efficiency of fault tolerance check pointing, which is potentially needed for shrinking CMOS processor logic that leads to less reliable chips. There is a major impact from this technology on software and computing. SCM provides orders of magnitude increase in capacity with near-DRAM latency which would push software towards in-memory computing.

Resistive Computing, Neuromorphic Computing and Quantum Computing are promising technologies that may be suitable for new hardware accelerators but less for new processor logic. Resistive computing promises a reduction in power consumption and massive parallelism. It could enforce datacentric and reconfigurable computing, leading away from the Von-Neumann architecture. Humans can easily outperform currently available high-performance computers in tasks like vision, auditory perception and sensory motor-control. As Neuromorphic Computing would be efficient in energy and space for artificial neural network applications, it would be a good match for these tasks. More lack of abilities of current computers can be found in the area of unsolved problems in computer science. Quantum Computing might solve some of these problems, with important implications for public-key cryptography, searching, and a number of specialized computing applications.

Summary of Potential Long-Term Impacts of Disruptive Technologies for HPC Software and Applications

New technologies will lead to new hardware structures with demands on system software and programming environment and also opportunities for new applications.

CMOS scaling will require system software to deal with higher fault rate and less reliability. Also programming environment and algorithms may be affected, e.g., leading to specifically adapted approximate computing algorithms.

The most obvious change will result from changes in memory technology. NVM will prevail independent of the specific memristor technology that will win. The envisioned Storage-Class Memory (SCM) will influence system software and programming environments in several ways:

- Memory and storage will be accessed in a uniform way.
- Computing will be memory-centric.
- Faster memory accesses by the combination of NVM and photonics will lead to a shallower memory hierarchy envisioning a flat memory where latency does not matter anymore.
- Read accesses will be faster than write accesses, though, software needs to deal with the read/write disparity, e.g., by database algorithms that favour more reads over writes.
- NVM will allow in-memory checkpointing, i.e. checkpoint replication with memory to memory operations.

- Software and hardware needs to deal with limited endurance of NVM memory.

A lot of open research questions arise from these changes for software.

Full 3D stacking may pose further requirements to system software and programming environments:

- The higher throughput and lower memory latency when stacking memory on top of processing may require changes in programming environments and application algorithms.
- Stacking specialized (e.g. analog) hardware on top of processing and memory elements lead to new (embedded) high-performance applications.
- Stacking hardware accelerators together with processing and memory elements require programming environment and algorithmic changes.
- 3D multicores require software optimizations able to efficiently utilize the characteristics of 3rd dimension, .i.e. e.g., different latencies and throughput for vertical versus horizontal interconnects.
- 3D stacking may to new form factors that allow for new (embedded) high-performance applications.

Photonics will be used to speed up all kind of interconnects – layer to layer, chip to chip, board to board, and compartment to compartment with impacts on system software, programming environments and applications such that:

- A flatter memory hierarchy will be reached (combined with 3D stacking and NVM) requiring software changes for efficiency redefining what is local in future.
- It is mentioned that energy-efficient Fourier-based computation is possible as proposed in the Optalysys project.
- The intrinsic end-to-end nature of an efficient optical channel will favour broadcast/multicast based communication and algorithms.
- A full photonic chip will totally change software in a currently rarely investigated manner.

A number of new technologies will lead to new accelerators. We envision programming environments that allow defining accelerator parts of an algorithm independent of the accelerator itself. OpenCL is such a language distinguishing “general purpose” computing parts and accelerator parts of an algorithm, where the accelerator part can be compiled to GPUs, FPGAs, or many-cores like the Xeon Phi. Such programming environment techniques and compilers have to be enhanced to improve performance portability and to deal with potentially new accelerators as, e.g., neuromorphic chips, quantum computers, in-memory resistive computing devices etc. System software has to deal with these new possibilities and map computing parts to the right accelerator.

Neuromorphic Computing is particularly attractive for applying artificial neural network and deep learning algorithms in those domains where, at present, humans outperform any currently available high-performance computer, e.g., in areas like vision, auditory perception, or sensory motor-control. Neural information processing is expected to have a wide applicability in areas that require a high degree of flexibility and the ability to operate in uncertain environments where information usually is partial, fuzzy, or even contradictory. The success of the IBM Watson computer is an example for such new application possibilities.

It is envisioned that neuromorphic computing could help understanding the multi-level structure and function of the brain and even reach an electronic replication of the human brain at least in some areas such as perception and vision.

Quantum Computing potentially solves problems impossible by classical computing, but posts challenges to compiler and runtime support. Moreover, quantum error correction is needed due to high error rates (10^{-3}). Applications of quantum computers could be new encryptions, quantum search, quantum random walk, etc.

Resistive Computing may lead to massive parallel computing based on data-centric and reconfigurable computing paradigms. In memory computing algorithms may be executed on specialised resistive computing accelerators.

Quantum Computing, Resistive Computing as well as Graphene and Nanotube-based computing are still highly speculative hardware technologies.

3. Sustaining Technology (improving HPC HW in ways that are generally expected)

Continuous CMOS scaling

Current (2016) high-performance multiprocessors feature 14 to 16nm technology. In April 2015, TSMC announced that the 10nm production would begin at the end of 2016. On 23 May 2015, Samsung Electronics showed off a 300mm wafer based on 10nm FinFET chips. Intel delayed their 10nm manufactured Cannonlake processor until the second half of 2017 [5] due to problems with the manufacturing process with 10nm technology. Intel's difficulties and changed plans show the continuing challenges with keeping pace with Moore's law.

Continuing Moore's Law and managing power and performance tradeoffs remain as the key drivers of the International Technology Roadmap For Semiconductors 2015 Edition (ITRS 2015) [1] grand challenges. Silicon scales according to the ITRS 2013 Roadmap until around 7 to 8nm in 2025 and 6 to 5nm in 2028 for MPUs or ASICs. DRAM half pitch (i.e., half the distance between identical features in an array) is projected to scale down to 10nm in 2025 and 7.7nm in 2028 allowing up to 32 GBits per chip. However, DRAM scaling below 20 nm is very challenging [1]. This results in an increasing cost of transistors at nodes below 10nm: the cost per transistor may increase from one technology node to the next [2].

The ITRS roadmap does not guarantee that silicon-based CMOS will extend that far because transistors with a gate length of 6 nm or smaller are significantly affected by quantum tunneling [3]. As a result of the limited further CMOS scaling the ITRS redirected their focus [4].

One trend to improve the density on chips will be 3D integration. A revolutionary DRAM/SRAM replacement will be needed [1]. As a result, non-silicon extensions of CMOS, using III-V materials or Carbon nanotube/nanowires, as well as non-CMOS platforms, including molecular electronics, spin-based computing, and single-electron devices, have been proposed [3].

Impact on hardware: "Scaling von Neumann systems leads to steadily increasing power consumption, high voltage density and high clock frequency leading away from the operating points of a biological brain" [3].

For a higher integration density, new materials and processes will be necessary. Since there is a lack of knowledge of the fabrication process of such new materials, the reliability might be lower, which may result in the need of integrated fault-tolerance mechanisms [1].

References

- [1] Semiconductor Industry Association. "International technology roadmap for semiconductors (ITRS), 2015 edition." Hsinchu, Taiwan, 2015.
- [2] HiPEAC Vision, 2015
- [3] en.wikipedia.org/wiki/10_nanometer
- [4] <http://www.nature.com/nnano/journal/v11/n2/full/nnano.2016.8.html>



EUROLAB-4-HPC

[5] <http://arstechnica.com/gadgets/2015/07/intel-confirms-tick-tock-shattering-kaby-lake-processor-as-moores-law-falters/>

Die Stacking and 3D-Chip

Die Stacking and 3D chip integration denote the concept of stacking integrated circuits (e.g. processors and memories) vertically in multiple layers. 3D packaging assembles vertically stacked dies in a package, e.g., system-in-package (SIP) and package-on-package (POP).

Die stacking can be achieved in a stacking approach by connecting separately manufactured wafers or dies vertically either via wafer-to-wafer, die-to-wafer, or even die-to-die integration. The mechanical and electrical contacts are realized either by wire bonding as in SIP and POP devices or microbumps. SIP is sometimes listed as a 3D stacking technology, although it should be better denoted as 2.5 D technology.

Another approach is arranging dies (called chiplets) horizontally connected with Interposers onto silicon substrate. The advantages of 3D technology based on Interposer are numerous: Firstly, short communication distance between dies, thus reducing communication load and then reducing communication power consumption. Secondly, the possibility of stacking dies from various heterogeneous technologies, like stacking memory on top of logic like flash, nonvolatile memories, or even photonic devices, in order to benefit of the best technology where it best fits. And thirdly, an improved system yield and cost by partitioning the system in a divide & conquer approach: multiple similar dies are fabricated, tested and sorted before the final 3D assembly, instead of fabricating ultra large dies with much reduced yield.

Die stacking can also be achieved by stacking active layers vertically on a single wafer in a monolithic approach. Such kind of 3D chip integration does not use off-chip signaling for communication but it applies direct signaling between layers. Contacts are implemented in true 3D technology without mechanical contacts using inductive or capacitive effects or by vertical conductive channels through the chip substrate, so-called through-silicon-vias (TSV).

Since the TSV technology offers the densest connectivity, it is currently the most promising and favored 3D stacking technology for future high-performance microprocessors. Besides, there is also monolithic 3D technology, where layers are grown on top of the other. This is also more compact, and allows smaller grain integration between layers.

Current state: The monolithic approach of die stacking is already used in 3D flash memories from Samsung and also for smart sensors. Commercial prototypes of 3D technology date back until 2004 when Tezzaron released a 3D IC microcontroller [1]. Intel evaluated chip stacking for a Pentium4 already in 2006 [2]. Recent multicore designs using Tezzaron's technology include the 64 core 3D-MAPS (3D MAssively Parallel processor with Stacked memory) research prototype from 2012 [3] [4] and the Centip3De with 64 ARM Cortex-M3 Cores also from 2012 [5]. Fabs are able to handle 3D packages (e.g. [6]). In 2011 IBM announced 3D chip production process [7]. Intel announced "3D XPoint" memory in 2015 (assumed to be 10x the capacity of DRAM and 1000x faster than NAND flash [8]). Both NVIDIA and AMD already exploit the high-bandwidth and low latencies given by 3D stacked memories for a high-dense memory processor, called high-bandwidth memory (HBM). AMD's GPUs based on the Fiji architecture with HBM are available since 2015, and NVIDIA released Pascal-based GPUs in 2016 [17]. A direction towards future 3D stacking of memory dies on

processor dies is the Hybrid Memory Cube from Micron. It stacks multiple DRAM dies and a separate layer for a controller which is vertically linked with the DRAM dies. This interposer approach is used in high end FPGAs to reduce cost.

Perspective: 3D NAND Flash may be prevailing. 3D flash memories may enable SSDs with up to 10 TB of capacity in the short term [9]. In 2007, earliest potential was seen in memory stacks for mobile applications [10]. It is to expect that 3D chip technology will widely enter the market for mainstream architectures within the next 5 years. Representative for this current development are, e.g., Intel's Xeon Phi Knights Landing processors which will be equipped with package-integrated DRAMs in 2016 as a result of their cooperation with Micron.

It is also to be expected that in a long-term perspective the technology will be expanded progressively from 3D packaging technologies towards real 3D chip stacking and possibly towards 3D ICs in 3D packages in order to profit from all the benefits such technology will offer in particular for HPC architectures.

The main challenge in establishing this 3D chip stacking technology is gaining control of the thermal problems that have to be overcome to realize reliably very dense 3D interconnections. This requires the availability of appropriate design tools, which are explicitly supporting 3D layouts. Both topics represent an important issue for research in the next 10 to 15 years.

Impact on hardware: 3D stacking has a series of beneficial impacts on the hardware in general and on the possibilities how to design future processor-memory-architectures in particular. Wafers can be partitioned into smaller dies because comparatively long horizontally running links are relocated to the third dimension and thus enable smaller form factors. 3D stacking also enables heterogeneity, by integrating layers, manufactured in different processes, e.g., different memory technologies, like SRAM, DRAM, Spin-transfer-torque RAM (STT-RAM) and also memristor technologies, which would be incompatible among each other in monolithic circuits. Due to short connection wires, reduction of power consumption is to be expected. Simultaneously, a high communication bandwidth between layers connected with TSVs can be expected leading to particularly high processor-to-memory bandwidth.

The last-level caches will probably be the first to be affected by 3D stacking technologies, which will increase bandwidth and reduce latencies by a large cache memory stacked on top of logic circuitry. In a further step it is consequent to expand 3D chip integration also to main memory in order to make a strong contribution in reducing decisively the current memory wall which is one of the strongest obstructions in getting more performance in HPC systems. Furthermore, possibly between 2026 and 2030, 3D arithmetic units will undergo the same changes ending up in complete 3D many-core microprocessors, which are optimized in power consumption due to reduced wire lengths.

3D stacking will also be used to scale flash memories, because 2D NAND flash technology does not scale beyond 16 nm [9, 12]. 3D stacking can also be used for image sensors. A technology was presented by Olympus in which more than 4 million microbumps have been used for stacking a 16 megapixel array sensor directly on top of a circuit implementing a global shutter control logic. Sony used TSV technology to

combine image sensors directly with column-parallel analogue-digital-converters and logic circuits [13,14].

References

- [1] Tezzaron 3D-IC Microcontroller Prototype [Online] February 11, 2016. http://www.tachyonsemi.com/OtherICs/3D-IC_8051_prototype.htm
- [2] Black, B.; Annavaram, M.; Brekelbaum, N.; DeVale, J.; Lei Jiang; Loh, G.H.; McCauley, D.; Morrow, P.; Nelson, D.W.; Pantuso, D.; Reed, P.; Rupley, J.; Sadasivan Shankar; Shen, J.; Webb, C., "Die Stacking (3D) Microarchitecture," in *International Symposium on Microarchitecture (MICRO)*, pp.469-479, 2006
- [3] <http://arch.ece.gatech.edu/research/3dmaps/3dmaps.html>
- [4] Dae Hyun Kim; Athikulwongse, K.; Healy, M.; Hossain, M.; Moongon Jung; Khorosh, I.; Kumar, G.; Young-Joon Lee; Lewis, D.; Tzu-Wei Lin; Chang Liu; Panth, S.; Pathak, M.; Minzhen Ren; Guan hao Shen; Taigon Song; Dong Hyuk Woo; Xin Zhao; Joung ho Kim; Ho Choi; Loh, G.; Hsien-Hsin Lee; Sung Kyu Lim, "3D-MAPS: 3D Massively parallel processor with stacked memory," in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2012 IEEE International*, pp.188-190, 2012
- [5] Fick, D.; Dreslinski, R.G.; Giridhar, B.; Gyouho Kim; Sangwon Seo; Fojtik, M.; Satpathy, S.; Yoonmyung Lee; Daeyeon Kim; Liu, N.; Wieckowski, M.; Chen, G.; Mudge, T.; Sylvester, D.; Blaauw, D., "Centip3De: A 3930DMIPS/W configurable near-threshold 3D stacked system with 64 ARM Cortex-M3 cores," in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC)*, pp.190-192, 19-23, 2012
- [6] 3D & Stacked-Die Packaging Technology Solutions [Online] February 11, 2016. <http://www.amkor.com/go/3D-Stacked-Die-Packaging>
- [7] IBM Press Release [Online], in German, February 11, 2016. <http://www-03.ibm.com/press/de/de/pressrelease/36129.wss>
- [8] Intel® Optane™: Supersonic memory revolution to take-off in 2016 [Online] February 11, 2016. <http://www.intel.eu/content/www/eu/en/it-managers/non-volatile-memory-idf.html>
- [9] Intel offers ingenious piece of 10TB 3D NAND chipper [Online] February 11, 2016. http://www.theregister.co.uk/2014/11/21/intel_offering_an_ingenuous_piece_of_10tb_3d_nand_chippery/
- [10] Lu , Jian-Qiang; Rose, Ken; Vitkavage, Susan, "3D Integration: Why, What, Who, When?" in *Future Fab International , Issue 23, 2007*, http://homepages.rpiscrows.us/~luj/FutureFab23_Luj_Reprint.pdf.
- [11] AMD's high-bandwidth memory explained [Online] February 11, 2016. <http://techreport.com/review/28294/amd-high-bandwidth-memory-explained>
- [12] Eun-Seok Choi; Hyun-Seung Yoo; Han-Soo Joo; Gyu-Seog Cho; Sung-Kye Park; Seok-Kiu Lee, "A Novel 3D Cell Array Architecture for Terra-Bit NAND Flash Memory," in *3rd IEEE International Memory Workshop (IMW)*, pp.1-4,2011
- [13] Kondo, T.; Takemoto, Y.; Kobayashi, K.; Tsukimura, M.; Takazawa, N.; Kato, H.; Suzuki, S.; Aoki, J.; Saito, H.; Gomi, Y.; Matsuda, S.; Tadaki, Y., "A 3D stacked CMOS image sensor with 16Mpixel global-shutter mode and 2Mpixel 10000fps mode using 4 million interconnections," in *Symposium on VLSI Circuits (VLSI Circuits), 2015*, pp.C90-C91, 2015
- [14] ISSCC 2013: Sony Stacked Sensor Presentation (ISSCC 2013) [Online] February 11, 2016 <http://image-sensors-world-blog.blogspot.de/2013/02/isscc-2013-sony-stacked-sensor.html>
- [15] Y. Xie, J. Zhao, *Die-stacking Architecture*, Morgan & Claypool Publishers series, Synthesis Lectures on Computer Architecture, 2016.



[16] Y. Xie, J. Cong, and S. Sapatnekar, Three-Dimensional Integrated Circuit Design: EDA, Design and Microarchitectures. Springer, 2009

[17] <https://images.nvidia.com/content/pdf/tesla/whitepaper/pascal-architecture-whitepaper.pdf>

4. Disruptive Technology in Hardware/VLSI (innovation that creates a new line of HPC hardware superseding existing HPC techniques)

Non-volatile Memory (NVM) Technologies

Currently NAND Flash is the most common NVM technology, which finds its usages on SSDs, memory cards and memory sticks. NAND flash uses floating-gate transistors for storing single bits. This technology is facing a big challenge, because scaling down decreases the endurance and performance significantly [25]. Hence the importance of other NVM technologies increases.

Resistive memories, i.e. memristors, are an emerging class of non-volatile memory technology. The memristors electrical resistance is not constant but depends on the history of current that had previously flowed through the device. The device remembers its history—the so-called non-volatility property: when the electric power supply is turned off, the memristor remembers its most recent resistance until it is turned on again [1].

Among the most prominent memristor candidates and close to commercialization are phase change memory (PCM) [2, 3, 4, 5, 6], metal oxide resistive random access memory (RRAM or ReRAM) [7, 8], and conductive bridge random access memory (CBRAM) [9].

PCM can be integrated in the CMOS process and the read/write latency is only by tens of nanoseconds slower than DRAM whose latency is roughly around 100ns. The write endurance is hundreds of millions of writes per cell at current processes. This is why PCM is currently positioned only as a Flash replacement. [21]. RRAM offers a simple cell structure which enables reduced processing costs. The endurance can be more than 50 million cycles and the switching energy is very low [22]. RRAM can deliver 100x lower read latency and 20x faster write performance compared to NAND Flash [23]. CBRAM can also write with relatively low and with high speed. The read/write latencies are close to DRAM.

Spintronics is the technology of manipulating the spin state of electrons. Instead of using the electrons charge, spin states can be utilized as a substitute in logical circuits or in traditional memory technologies like SRAM. An STT-RAM [10] memory cell stores data in a magnetic tunnel junction (MTJ). Each MTJ is composed of two ferromagnetic layers (free and reference layers) and one tunnel barrier layer (MgO). If the magnetization direction of the magnetic fix reference layer and the switchable free layer is anti-parallel, resp. parallel, a high, resp. a low, resistance is adjusted, representing a digital "0" or "1". Recently it was reported that by adjusting intermediate magnetization angles in the free layer 16 different states can be stored in one physical cell, enabling to realize multi-cell storages in MTJ technology [11].

The read latency and read energy of STT-RAM is expected to be comparable to that of SRAM. The expected 3x higher density and 7x less leakage power consumption in the STT-RAM makes it suitable for replacing SRAMs to build large NVMs. However, a write operation in an STT-RAM memory consumes 8x more energy and exhibits a 6x longer latency than a SRAM. Therefore, minimizing the impact of inefficient writes is critical for successful applications of STT-RAM [12].

NRAM, short for Nano Ram is a proprietary technology of Nantero. The ram uses a fabric of carbon nanotubes (CNT) for saving bits. The resistive state of the CNT fabric determines, whether a one or a zero is saved in a memory cell. The resistance depends on if the CNTs are in contact with each other. With the help of a small voltage, the CNTs can be brought into contact or be separated. Reading out a bit means to measure the resistance. Nantero claims that their technology features the same read- and write latencies as DRAM, has a high endurance and reliability even in high temperature environments and is low power with essentially zero power consumption in standby mode. Furthermore NRAM is compatible with existing CMOS fabs without needing any new tools or processes, and it is scalable even to below 5nm [26].

Current state: IBM announced MLC-PCM technology replacing flash. Intel and Micron announced the new Breakthrough Memory 3D XPoint Technology [14] as revolutionary flash replacement. It is expected that the X-Point technology could become the dominating technology as an alternative to RAM devices offering in addition NVM property in the next ten years.

IBM also developed a neuromorphic core with a 64-K-PCM-cell as Synaptic-Array (256 Axone x 256 Dendrite) to implement SNNs (Spiking Neural Networks) [19].

Adesto is currently offering CBRAM technology in their serial memory chips [24].

The circuit-level performance, energy, and area model of the emerging non-volatile memory simulator NVSim [20] allows the investigation of architectural structures for future NVM based high-performance computers.

Perspective: It is foreseeable, that other NVM technologies will supersede current flash memory. PCM for instance might be 1000 times faster and 1000 times more resilient. Some NVM technologies have been considered as a feasible replacement for SRAM [15, 16, 17]. Studies suggest that replacing SRAM with STT-RAM could save 60% of LLC energy with less than 2% performance degradation [15].

It is unclear when most of the new technologies may be mature enough and which of them will prevail. But this is not of importance, because all have the same goal, namely to revolutionize the current storage technology.

Impact on hardware: Memristors will deliver non-volatile memory which can be used potentially in addition to DRAM, or as a complete replacement.

The latter will lead to a new Storage-Class Memory (SCM), i.e., a technology that blurs the distinction between memory and storage by enabling new data access modes and protocols that serve both 'memory' and 'storage'. These new SCM types of non-volatile memory could be integrated on-chip with the microprocessor cores as they use CMOS-compatible sets of materials and require different device fabrication techniques than flash. In a VLSI post-processing step they can be integrated on top of the last metal layer (see the note on Back-end of line service in section Resistive Computing). One of the challenges for the next decade is the provision of appropriate interfacing circuits between the SCMs and the microprocessor cores. The benefits of memristor devices in integration density, energy consumption and access times may not get lost by costly interface circuitry. This holds in particular for exploiting the multi-level cell storage capability of NVMs for future systems, e.g., for big data applications. Moreover,

memristors offer orders of magnitude faster read/write accesses and also much higher endurance. They are resistive switching memory technologies, and thus rely on different physics than that of storing charge on a capacitor as is the case for SRAM, DRAM and Flash [18].

Spin-transfer torque magnetic random access memory (STT-RAM) devices are also an important class of non-volatile memory that primarily targets the replacement of DRAM, e.g., in Last-Level Caches (LLC). However, the asymmetric read/write energy and latency of NVM technologies introduces new challenges in designing memory hierarchies. Spintronic allows integration of logic and storage at lower power consumption.

Also new hybrid PCM / Flash SSD chips could emerge with a processor-internal last-level cache (STT-RAM), main processor memory (PCRAM), and storage class memory (ReRAM) [18].

References

- [1] en.wikipedia.org/wiki/Memristor
- [2] B. C. Lee, P. Zhou, J. Yang, Y. Zhang, B. Zhao, E. Ipek, O. Mutlu, and D. Burger, "Phase-change technology and the future of main memory," *IEEE Micro*, vol. 30, pp. 143–143, Jan. 2010.
- [3] C. Lam, "Cell design considerations for phase change memory as a universal memory," in *VLSI Technology, Systems and Applications*, pp. 132–133, 2008.
- [4] B. C. Lee, E. Ipek, O. Mutlu, and D. Burger, "Architecting phase change memory as a scalable dram alternative," in *36th Annual International Symposium on Computer Architecture (ISCA-2009)*, pp. 2–13, 2009.
- [5] M. K. Qureshi, V. Srinivasan, and J. A. Rivers, "Scalable high performance main memory system using phase-change memory technology," in *36th Annual International Symposium on Computer Architecture (ISCA-2009)*, pp. 24–33, 2009.
- [6] P. Zhou, B. Zhao, J. Yang, and Y. Zhang, "A Durable and Energy Efficient Main Memory Using Phase Change Memory Technology," in *36th Annual International Symposium on Computer Architecture (ISCA-2009)*, pp. 14-23, 2009.
- [7] C. Xu, D. Niu, N. Muralimanohar, R. Balasubramonian, T. Zhang, S. Yu, and Y. Xie, "Overcoming the challenges of crossbar resistive memory architectures," in *IEEE 21st International Symposium on High Performance Computer Architecture (HPCA 2015)*, pp. 476–488, Feb 2015.
- [8] C. Xu, X. Dong, N. Jouppi, and Y. Xie, "Design implications of memristor-based rram cross-point structures," in *Design, Automation Test in Europe Conference Exhibition (DATE), 2011*, pp. 1–6, March 2011.
- [9] William Wong: *Conductive Bridging RAM*. electronic design, 2014, <http://electronicdesign.com/memory/conductive-bridging-ram>
- [10] D. Apalkov, A. Khvalkovskiy, S. Watts, V. Nikitin, X. Tang, D. Lottis, K. Moon, X. Luo, E. Chen, A. Ong, A. Driskill-Smith, and M. Krounbi, "Spin-transfer Torque Magnetic Random Access Memory (STT-MRAM)," *Journal Emerg. Technol. Comput. Syst.*, vol. 9, pp. 13:1–13:35, May 2013.
- [11] D. Bernard, "Spintronic devices for memristor applications, Talk at Meeting of EU COST ACTION MemoCis IC1401, "Memristors: at the crossroad of Devices and Applications", Milano, 28th March 2016.

- [12] H. Noguchi, K. Kushida, K. Ikegami, K. Abe, E. Kitagawa, S. Kashiwada, C. Kamata, A. Kawasumi, H. Hara, and S. Fujita, "A 250-MHz 256b-I/O 1-Mb STT-MRAM with advanced perpendicular MTJ based dual cell for nonvolatile magnetic caches to reduce active power of processors," in 2013 Symposium on VLSI Technology (VLSIT), pp. 108–109, 2013.
- [13] Gary Hilson: IBM Tackles Phase-Change Memory Drift, Resistance. EETimes 5/1/2015. http://www.eetimes.com/document.asp?doc_id=1326477
- [14] <http://www.intel.com/content/www/us/en/architecture-and-technology/non-volatile-memory.html>
- [15] H. Noguchi, K. Ikegami, N. Shimomura, T. Tetsufumi, J. Ito, and S. Fujita, "Highly reliable and low-power nonvolatile cache memory with advanced perpendicular STT-MRAM for high-performance CPU," in Symposium on VLSI Circuits Digest of Technical Papers, pp. 1–2, June 2014.
- [16] H. Noguchi, K. Ikegami, K. Kushida, K. Abe, S. Itai, S. Takaya, N. Shimomura, J. Ito, A. Kawasumi, H. Hara, and S. Fujita, "A 3.3ns-access-time 71.2 uW/MHz 1Mb embedded STT-MRAM using physically eliminated read-disturb scheme and normally-off memory architecture," in IEEE International Solid-State Circuits Conference (ISSCC), pp. 1–3, Feb 2015.
- [17] J. Ahn, S. Yoo, and K. Choi, "DASCA: Dead Write Prediction Assisted STT-RAM Cache Architecture," in Proceedings of the 2014 IEEE 20th International Symposium on High Performance Computer Architecture, HPCA '14, 2014.
- [18] Evangelos Eleftheriou, Future Non-Volatile Memories: Technology Trends and Applications, Keynote, CSW Milan, 2015
- [19] Neuromorphes System für SNNs, 09.12.2015; <http://www.elektroniknet.de/halbleiter/prozessoren/artikel/126062/>
- [20] X. Dong, C. Xu, Y. Xie, and N. Jouppi, "NVSim: A Circuit-Level Performance, Energy, and Area Model for Emerging Nonvolatile Memory," IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 31, pp. 994–1007, July 2012.
- [21] Lee, B. C., Ipek, E., Mutlu, O., & Burger, D. (2009). Architecting phase change memory as a scalable dram alternative. ACM SIGARCH Computer Architecture News, 37(3), 2-13.
- [22] Govoreanu, B., Kar, G. S., Chen, Y. Y., Paraschiv, V., Kubicek, S., Fantini, A., ... & Jossart, N. (2011, December). 10× 10nm² Hf/HfO_x crossbar resistive RAM with excellent performance, reliability and low-energy operation. In Electron Devices Meeting (IEDM), 2011 IEEE International (pp. 31-6). IEEE.
- [23] <http://www.crossbar-inc.com/technology/rram-advantages/>
- [24] <http://www.adestotech.com/products/mavriq/>
- [25] <http://www.anandtech.com/show/7237/samsungs-vnand-hitting-the-reset-button-on-nand-scaling>
- [26] <http://nantero.com/technology/>

Photonics

The general idea is to replace electrons with photons in intra-chip connections, inter-chip, memory and logic. First, interconnects between devices and chips are built from optical fiber, which allows higher throughput and possibly lower latency (head-flit latency). Such connections already exist (e. g. Thunderbolt or optical PCI Express by Intel). Optical inter-chip signals are then expected to be conveyed also on different mediums to facilitate integrability with CMOS process, e.g., polycarbonate as in some IBM research prototypes and commercial solutions.

The next step is to use optical interconnects to connect chips on the same circuit board. For future devices, also intra-chip connections may be optical. In general, the trend is going towards photonics integrated with electronics and closer to the chips and cores. However, conversion from photons to electrons is costly and for this reason there are currently strong efforts in improving the crucial physical modules of an integrated optical channel (e.g. modulators, photodetectors and thermally stable and efficiently integrated laser sources). Therefore, overall improvement in the effective adoption of photonics links closer and closer to the cores is to be expected.

On another direction, to eliminate the electrical-optical conversion overhead, research tries to build chips which completely are based on photonics. Silicon photonics describes the attempt to integrate photonics in CMOS-like production [1].

Optical or photonic computing uses photons produced by lasers or diodes for computation. Most research projects focus on replacing current computer components with optical equivalents, resulting in an optical digital computer system processing binary data. This approach appears to offer the best short-term prospects for commercial optical computing, since optical components could be integrated into traditional computers to produce an optical-electronic hybrid [2]. This approach is currently far from general but it is however suitable for some specific application domains, e.g., extremely energy-efficient Fourier-based computation proposed in the Optalysys project (<http://optalysys.com>) [5].

Current State: Optical fiber connections between devices exist. Currently, optical PCI Express connections allow high bandwidth and high speed over long distances (up to around 80m). Dedicated units for conversion between optical and electronic signals are required. Some integrated photonics solutions exist and are mainly aimed at replacing point-to-point electric wires (e.g. IBM HPC systems).

Perspective: Inter-chip connections on a single circuit board will become available soon. This requires silicon photonics, which allow the integration of photonics in a CMOS-like manufacturing process. Silicon photonics implement lasers, detectors, and waveguides on-chip with silicon only [3].

Research is working on optical intra-chip and inter-chip connections, with prototypes being already available. In this direction researchers have already identified the importance of a vertical design space exploration and design of a computer system endowed with integrated photonics. This can be combined with 3D technologies, replacing an active silicon interposer by a photonic interposer. Further challenges will arise from the evidence in current research proposals and prototypes that lower-layer design choices (e.g. physical layer, topologies, access strategies, sharing of resources),

can have a significant impact in higher layers of the design (e.g. NoC-wise and up to memory coherence and programming model implications) and vice versa. This is mainly due to the scarce experience in using photonics technology for serving computing needs (close to processing cores requirements) and, most of all, due to the intrinsic end-to-end nature of an efficient optical channel, which is completely opposed to the well-established and mature knowledge of “store-and-forward” electronic communication paradigm. Then, intrinsic low-latency properties of optical interconnection (on-chip and inter-chip) could imply a re-definition of what is local in a future computing system, especially at high-scale like in a perspective HPC, together with the programming paradigms able to take advantage of the induced new optimal organization of the overall machine.

Further research targets photonic non-volatile memory [4]. This could reduce latencies of memory accesses by eliminating costly optoelectronic conversions. A revolution of micro architecture design is possible, since latencies and differences in speed between CPU and main memory in fully optical chips will not exist anymore.

There are disagreements between researchers about the future capabilities of optical computers: Will they be able to compete with semiconductor-based electronic computers on speed, power consumption, cost, and size? For optical logic to be competitive beyond a few niche applications, major breakthroughs in non-linear optical device technology would be required, or perhaps a change in the nature of computing itself [2].

The English company Optalysys, however, is of different opinion and announces that “Optalysys’s initial products will launch in 2017 and are expected to enable existing computers to achieve HPC-levels of performance up to an equivalent processing rate of 9 Petaflops – comparable to the 5th fastest computer in the world today. Following that we plan to pursue the design of larger systems capable of achieving multiple Exaflops by 2020” [5].

Impact on hardware: With both memory and connections becoming faster, 3rd level caches in Von Neumann architectures may become obsolete. Lots of main memory could be accessed with small latencies. Possibly, if the whole microarchitecture is implemented on silicon photonics, computational units and memory could work with the same speed. The elimination of the von-Neumann bottleneck promises completely new and different architectures.

References

- [1] http://researcher.watson.ibm.com/researcher/view_group.php?id=2757
- [2] http://en.wikipedia.org/wiki/Optical_computing
- [3] <http://adsabs.harvard.edu/abs/2005Natur.433..292R>
- [4] <http://www.nature.com/nphoton/journal/vaop/ncurrent/full/nphoton.2015.182.html>
- [5] <http://optalysys.com/optalysys-prototype-proves-optical-processing-technology-will-revolutionise-big-data-analysis-computational-fluid-dynamics-cfd/>

5. Disruptive Technology (alternative ways of computing)

Resistive Computing

Apart from using memristors as non-volatile memory, there are several other ways to use memristors in computing [1, 2]. Using memristors as memristive synapses in neuromorphic computing [2, 3, 4] and using memristors in quantum computing [2] are discussed in separate sections. In this section, resistive computing is discussed.

In resistive computing, logic circuits are built by memristors [5]. Memristive gates have a lower leakage power, but switching is slower than in CMOS gates [2]. However, the integration of memory into logic allows to reprogram the logic, providing low power reconfigurable components [12] and can reduce energy and area constraints in principle due to the possibility of computing and storing in the same device (computing in memory). Memristors can also be arranged in parallel networks to enable massively parallel computing [13].

Resistive computing is one of the emerging and promising computing paradigms [5,6,7]. It takes the data-centric computing concept much further by interweaving the processing units and the memory in the same physical location using non-volatile technology, therefore significantly reducing not only the power consumption but also the memory bottleneck. Resistive devices such as memristors have been shown to be able to perform both storage and logic functions [5,8,9,10,11].

Resistive computing provides a huge potential as compared with the current state-of-the-art:

- It significantly reduces the memory bottleneck as it interweaves the storage, computing units and the communication [5,6,7].
- It features low power leakage [2].
- It enables maximum parallelism [7,13].
- It allows full configurability and flexibility [12].
- It provides order of magnitude improvements for the energy-delay product per operations, the computation efficiency, and performance per area [7].

Serial and parallel connections of memristors were proposed for the realization of Boolean logic gates with memristors by the so-called memristor ratio logic. In such circuits the ratio of the stored resistances in memristor devices is exploited for the set-up of Boolean logic. Memristive circuits realizing AND, OR gates and the implication function were presented in [14,15,19]. Hybrid memristive computing circuits consist of memristors and CMOS gates. The research of Singh [17], Xia et.al. [18], and Rothenbuhler et.al.[19] are representative for numerous proposals of hybrid memristive circuits, in which most of the Boolean logic operators are handled in the memristors and the CMOS transistors are mainly used for level restoration to retain defined digital signals.

Perspective: Resistive computing, if successful, will be able to significantly reduce the power consumption and enable massive parallelism; hence, increase computing energy and area efficiency by orders of magnitudes. This will transform computer systems into new highly parallel architectures and associated technologies, and enables the

computation of currently infeasible big data and data-intensive applications, fueling important societal changes.

Research on resistive computing is still in its infancy stage, and the challenges are substantial at all levels, including material and technology, circuit and architecture, tools and compilers, and algorithms. As of today most of the work is based on simulations and small circuit designs. It is still unclear when the technology will be mature and available. Nevertheless, some start-ups on memristor technologies are emerging such as KNOWM.

A couple of start-up companies appeared in 2015 on the market who offer memristor technology as BEOL (Back-end of line) service in which memristive elements are post-processed in CMOS chips directly on top of the last metal layers. Also some European institutes reported just recently at a workshop meeting “Memristors: at the crossroad of Devices and Applications” of the EU cost action 1401 MemoCis the possibility BEOL integration of their memristive technology to allow experiments with such technologies. This offers new perspectives in form of hybrid CMOS/memristor logic which use memristor networks for high-dense resistive logic circuits and CMOS inverters for signal restoration to compensate the loss of full voltage levels in memristive networks. Multi-level cell capability of memristive elements can be used to face the challenge to handle the expected huge amount of Zettabytes produced annually in a couple of years. Besides, proposals exist to exploit the multi-level cell storing property for ternary carry-free arithmetic [20], [21] or both compact storing of keys and matching operations in future associative memories realized with memristors [22].

Impact on hardware: Due to its nature which is Non-Von Neumann, resistive computing will significantly change the way we used to design our computers, both from software as well as from hardware perspective. It will enforce datacentric and reconfigurable computing. Hybrid memristive networks can reduce energy and area requirement of logic circuits compared to pure CMOS. Massively parallel networks of memristors could form specialized accelerators to solve NP-hard problems [7].

Currently the interfacing and the peripheral to access memristive elements costs a lot of energy. Appropriate low energy driver circuits and new design flows are necessary to face the challenges for memristive circuits to use them. At the end extreme low-energy consumption devices can be realized both for big data applications with computing-in-memory and high-performance embedded computing devices that can be operated completely by energy harvesting mechanisms thanks to the use of memristors which offer low energy consumption circuits, high storage densities and low number of latency states in arithmetic circuits due to carry-free additions with ternary number representations.

References

- [1] Massimiliano Di Ventra, Yuriy V. Pershin: The parallel approach, *Nature Physics* 9, 200–202 (2013)
- [2] Pershin, Y.V., Di Ventra, M.: Neuromorphic, Digital, and Quantum Computation With Memory Circuit Elements, *Proc. IEEE* , 100(6) 2071-2080 (2011)
- [3] Matthew D. Pickett, Gilberto Medeiros-Ribeiro, Stanley Williams: A scalable neuristor built with Mott memristors, *Nature Materials* 12, 114–117 (2013)



- [4] Sung Hyun Jo, Ting Chang, Idongesit Ebong, Bhavitavya B. Bhadviya, Pinaki Mazumder and Wei Lu: Nanoscale Memristor Device as Synapse in Neuromorphic Systems, *Nano Lett.*, 10 (4), pp 1297–1301 (2010)
- [5] Julien Borghetti, Gregory S. Snider, Philip J. Kuekes, J. Joshua Yang, Duncan R. Stewart, R. Stanley Williams: 'Memristive' switches enable 'stateful' logic operations via material implication, *Nature* 464, 873–876 (2010)
- [6] M. Di Ventra et al., "Memcomputing: a computing paradigm to store and process information on the same physical platform," arXiv preprint arXiv:1211.4487, 2012.
- [7] Said Hamdioui, Lei Xie, Hoang Anh Du Nguyen, Mottaqiallah Taouil, Koen Bertels, Henk Corporaal, Hailong Jiao, Francky Catthoor, Dirk Wouters, Linn Eike, Jan van Lunteren: "Memristor Based Computation-in-Memory Architecture for Data-Intensive Applications", *Proceedings of the 2015 Design, Automation & Test in Europe & Exhibition*, pp. 1718-1725, 2015.
- [8] G. Snider, "Computing with hysteretic resistor crossbars," *Applied Physics A*, vol. 80, pp. 1165–1172, 2005.
- [9] J. Borghetti et al., "Memristive switches enable stateful logic operations via material implication," *Nature*, vol. 464, pp. 873–876, 2010.
- [10] L. Gao et al., "Programmable cmos/memristor threshold logic," *IEEE Transactions on Nanotechnology*, vol. 12, pp. 115–119, 2013.
- [11] Lei Xie, Hoang Anh Du Nguyen, Mottaqiallah Taouil, Koen Bertels, Said Hamdioui: Fast boolean logic mapped on memristor crossbar. *The IEEE International Conference on Computer Design*, pp. 335-342, 2015.
- [12] Julien Borghetti, et al.: A hybrid nanomemristor/transistor logic circuit capable of self-programming, *Proc. Natl. Acad. Sci. USA* 106(6) 1699–1703 (2009)
- [13] Yuriy V. Pershin, Massimiliano Di Ventra: Solving mazes with memristors: A massively parallel approach, *Phys. Rev. E* 84, 046703 (2011)
- [14] J. J. Yang, D. B. Strukov, and D. R. Stewart. Memristive devices for computing. *Nature nanotechnology*, 8(1):13-24, 2013.
- [15] S. Kvatinisky, A. Kolodny, U. C. Weiser, and E. G. Friedman. Memristor-based imply logic design procedure. In *Proceedings of the 2011 IEEE 29th International Conference on Computer Design, ICCD'11*, pages 142-147, Washington, DC, USA, 2011. IEEE Computer Society.
- [16] S. Kvatinisky, N. Wald, G. Satat, A. Kolodny, U. Weiser, and E. Friedman. Mrl - memristor ratioed logic. In *Cellular Nanoscale Networks and Their Applications (CNNA)*, 2012 13th International Workshop on, pages 1-6, Aug 2012.
- [17] T. Singh. Hybrid memristor-cmos (memos) based logic gates and adder circuits. *CoRR*, abs/1506.06735, 2015.
- [18] Q. Xia, W. Robinett, M. W. Cumbie, N. Banerjee, T. J. Cardinali, J. J. Yang, W. Wu, X. Li, W. M. Tong, D. B. Strukov, and Others. Memristor- CMOS Hybrid Integrated Circuits for Reconfigurable Logic. *Nano letters*, 9(10):3640–3645, 2009.
- [19] A. Rothenbuhler, T. Tran, E. H. B. Smith, V. Saxena, and K. A. Campbell. Recon_gurable threshold logic gates using memristive devices. *Journal of Low Power Electronics and Applications*, 3(2):174-193, 2013.
- [20] A. El-Slehdar, A. Fouad, and A. G. Radwan. Memristor based n-bits redundant binary adder. *Microelectronics Journal*, 46(3):207-213, 2015.
- [21] D. Fey. Using the multi-bit feature of memristors for register files in signed-digit arithmetic units. *Semiconductor Science and Technology*, 29(10):104008, 2014.



[22] P. Junsangri, F. Lombardi, and J. Han. A memristor-based tcam (ternary content addressable memory) cell. In Nanoscale Architectures (NANOARCH), 2014 IEEE/ACM International Symposium on, pages 1-6, July 2014.

Neuromorphic Computing

Neuromorphic Computing, as developed by Carver Mead in the late 1980s, describes the use of very-large-scale integration (VLSI) systems containing electronic analog circuits to mimic neuro-biological architectures present in the nervous system.

The basic idea of Neuromorphic Computing is to exploit the massive parallelism of such circuits and to create low-power and fault-tolerant information-processing systems. Aiming at overcoming the big challenges of deep-submicron CMOS technology (power wall, reliability, and design complexity), bio-inspiration offers alternative ways to (embedded) artificial intelligence. The challenge is to understand, design, build, and use new architectures for nanoelectronic systems, which unify the best of brain-inspired information processing concepts and of nanotechnology hardware, including both algorithms and architectures [9]. A key focus area in further scaling and improving of cognitive systems is decreasing the power density and power consumption, and overcoming the CPU/memory bottleneck of conventional computational architectures [14].

In recent times, the term neuromorphic has also been used to describe analog, digital, and mixed-mode analog/digital VLSI and software systems that implement models of neural systems (for perception, motor control, or multisensory integration). The implementation of neuromorphic computing on the hardware level can be realized by oxide-based memristors, threshold switches and transistors [1, 2, 3, 4]. Such kind of research is still in its infancy.

Current state: Large scale neuromorphic chips exist based on CMOS technology, replacing processor cores by artificial neural networks. Research projects on neuromorphic computing are the following.

Mapping brain-like structures and processes into electronic substrates has recently seen a revival with the availability of deep-submicron CMOS technology. Large programs on brain-like electronic systems have been launched worldwide. At present, the largest programs are the SyNAPSE program (Systems of Neuromorphic Adaptive Plastic Scalable Electronics) in the US (launched in 2009, [10]) and the EC flagship Human Brain Project (launched in 2013, [11]).

SyNAPSE is a DARPA-funded program to develop electronic neuromorphic machine technology that scales to biological levels. More simply stated it is an attempt to build a new kind of computer with similar form and function to the mammalian brain. Such artificial brains would be used to build robots whose intelligence matches that of mice and cats. The ultimate aim is to build an electronic microprocessor system that matches a mammalian brain in function, size, and power consumption. It should recreate 10 billion neurons, 100 trillion synapses, consume one kilowatt (same as a small electric heater), and occupy less than two litres of space ([10]).

The “Cognitive Computing via Synaptronics and Supercomputing” (C2S2) project is a funded project from DARPA’s SyNAPSE initiative. Headed by IBM the group will turn to digital special-purpose hardware for brain emulation. The True North chip is an impressive outcome of this project integrating a two-dimensional on-chip network of 4096 digital application-specific cores (64 x 64) and over 400 Mio. bits of local on-chip memory (~100 Kb SRAM per core) to store synapses and neuron parameters as well as

256 Mio. individually programmable synapses on-chip. One million individually programmable neurons can be simulated time-multiplexed per chip, sixteen-times more than the current largest neuromorphic chip. The chip with about 5.4 billion transistors is fabricated in a 28nm CMOS process (4.3 cm² die size, 240µm x 390 µm per core). By device count, True North is the largest IBM chip ever fabricated and the second largest (CMOS) chip in the world. The total power, while running a typical recurrent network at biological real-time, is about 70mW resulting in a power density of about 20mW/cm² (about 26pJ) which is in turn comparable to the cortex but three to four orders-of magnitude lower compared to 50-100W/cm² for a conventional CPU [12].

Another US initiative is the Brain Corporation (qualcomm venture). It is a pioneer in developing novel algorithms based on the functioning of the nervous system, with applications to vision, motor control, and autonomous navigation. It is working with partners to design specialized hardware that will bring to market the next generation of smart consumer products with artificial nervous systems [5].

The Human Brain Project (HBP) is a European Commission Future and Emerging Technologies Flagship. The HBP aims to put in place a cutting-edge, ICT-based scientific research infrastructure that will allow scientific and industrial researchers to advance our knowledge in the fields of neuroscience, computing and brain-related medicine. The Project promotes collaboration across the globe, and is committed to driving forward European industry. Within the HBP the subproject SP9 designs, implements and operate a Neuromorphic Computing Platform with configurable Neuromorphic Computing Systems (NCS). The platform provides NCS based on physical (analogue or mixed-signal) emulations of brain models, running in accelerated mode (NM-PM1, wafer-scale implementation with about 200.000 analogue neurons on a wafer in 180nm CMOS), numerical models running in real time on digital multicore architectures (NM-MC1 with 18 ARM cores per chip in 130nm CMOS), and the software tools necessary to design, configure and measure the performance of these systems. The platform will be tightly integrated with the High Performance Analytics and Computing Platform, which will provide essential services for mapping and routing circuits to neuromorphic substrates, benchmarking and simulation-based verification of hardware specifications [12].

Closely related to HBP are the Blue Brain Project and the BrainScales project. The goal of the Blue Brain Project (EPFL and IBM, launched 2005): "... is to build biologically detailed digital reconstructions and simulations of the rodent, and ultimately the human brain. The supercomputer-based reconstructions and simulations built by the project offer a radically new approach for understanding the multilevel structure and function of the brain." The project uses an IBM Blue Gene supercomputer (100 TFLOPS, 10TB) with currently 8,000 CPUs to simulate ANNs (at ion-channel level) in software [15].

The European funded research project BrainScaleS (Brain-inspired multiscale computation in neuromorphic hybrid systems) aimed at understanding and emulating functions and interactions of multiple spatial and temporal scales in brain-information processing. Both, numerical simulations on Petaflop supercomputers and fundamentally different non Von Neumann hardware architectures were employed for this purpose. Within its broad scope of advancing neuromorphic computing, the

hardware part is a very-large-scale, mixed-signal implementation of a highly connected, adaptive network of analogue neurons. The basic element is the HICANN (High Input Count Analog Neural Network) chip hosting one analogue network core and necessary support circuitry for communication as well as controlling. The HICANN was implemented in a 180 nm CMOS technology, has a total of 112K synapses and 512 neuron circuits and is the basic component of the HBP NM-PM1 platform [13].

According to Olivier Temam's website [6], the following things have been achieved for hardware neural networks:

- ASIC-like energy efficiency on a digital CMOS and an analog design.
- Tolerance to permanent faults on both GPUs and a custom design.
- Tolerance to transient faults.
- That about half of the PARSEC benchmarks could benefit from an NN accelerator.
- A small-footprint high-throughput accelerator for enabling state-of-the-art machine-learning in data centers or embedded systems [7].
- Tape out of a 3D stacked NN to outline that 3D stacking might be a particularly suitable scalability path for neuromorphic architectures [8].

Perspective: Software implemented artificial neural networks on HPC-clusters, multi-cores (OpenCV), and GPGPUs (NVIDIA cuDNN) are already commercially used. FPGA acceleration of neural networks is available as well. From a short term perspective these software implemented neural networks may be accelerated by commercial transistor-based neuromorphic chips or accelerators. Future emerging hardware technologies, like memcomputing and 3D stacking [8] may bring neuromorphic computing to a new level and overcome some of the restriction of Von Neumann based VLSI systems in terms of scalability, power consumption or performance.

The building blocks for ICs and for the Brain are the same at nanoscale level: electrons, atoms, and molecules, but their evolutions have been radically different. The fact that reliability, low-power, reconfigurability, as well as asynchronicity have been brought up so many times in recent conferences and articles, makes it compelling that the Brain should be an inspiration at many different levels, suggesting that future nano-architectures could be neural-inspired. The fascination associated with an electronic replication of the human brain has grown with the persistent exponential progress of chip technology. The present decade 2010–2020 has also made the electronic implementation more feasible, because electronic circuits now perform synaptic operations such as multiplication and signal communication at energy levels of 10 fJ, comparable to biological synapses. Nevertheless, an all-out assembly of 10¹⁴ synapses will remain a matter of a few exploratory systems for the next two decades because of several challenges [9].

Impact on hardware: Neuromorphic computing would be efficient in energy and space and applicable as hardware accelerator.

Particularly attractive is the application of ANNs in those domains where, at present, humans outperform any currently available high-performance computer, e.g., in areas like vision, auditory perception, or sensory motor-control. Neural information processing is expected to have a wide applicability in areas that require a high degree of flexibility and the ability to operate in uncertain environments where information usually is partial, fuzzy, or even contradictory. Even more computational power may be

obtained by emerging technologies like quantum computing, molecular electronics, or novel nano-scale devices (memristor, spintronics, nanotubes (CMOL)) [9].

References

- [1] https://en.wikipedia.org/wiki/Neuromorphic_engineering
- [2] Pershin, Y.V., Di Ventra, M.: Neuromorphic, Digital, and Quantum Computation With Memory Circuit Elements, Proc. IEEE , 100(6) 2071-2080 (2011)
- [3] Matthew D. Pickett, Gilberto Medeiros-Ribeiro, Stanley Williams: A scalable neuristor built with Mott memristors, Nature Materials 12, 114–117 (2013)
- [4] Sung Hyun Jo, Ting Chang, Idongesit Ebong, Bhavitavya B. Bhadviya, Pinaki Mazumder and Wei Lu: Nanoscale Memristor Device as Synapse in Neuromorphic Systems, Nano Lett., 10 (4), pp 1297–1301 (2010)
- [5] <http://www.fundingpost.com/venturefund/venture-fund-profile.asp?fund=1012>
- [6] <http://pages.saclay.inria.fr/olivier.temam/homepage/research.html>
- [7] Tianshi Chen, Zidong Du, Ninghui Sun, Jia Wang, Chengyong Wu, Yunji Chen, and Olivier Temam. 2014. DianNao: a small-footprint high-throughput accelerator for ubiquitous machine-learning. ASPLOS '14. ACM, New York, NY, USA, 269-284.
- [8] Belhadj, B.; Valentian, A.; Vivet, P.; Duranton, M.; He, L.; Temam, O., "The improbable but highly appropriate marriage of 3D stacking and neuromorphic accelerators," CASES 2014, pp.1-9, 2014.
- [9] Rueckert, U., "Brain-Inspired Architectures for Nanoelectronics",. In: Hoefflinger B. (Ed.) "CHIPS2020, Vol. 2", Chap. 18, Springer, 2016.
- [10] <http://www.artificialbrains.com/darpa-synapse-program>
- [11] <http://www.humanbrainproject.eu>
- [12] Merolla, P.A. et al.: A million spiking-neuron integrated circuit with a scalable communication network and interface, Science 345, pp. 668-673, 2014.
- [13] <http://brainscales.kip.uni-heidelberg.de>
- [14] Evangelos Eleftheriou, Future Non-Volatile Memories: Technology Trends and Applications, Keynote, CSW Milan, 2015
- [15] <http://bluebrain.epfl.ch/page-56882-en.html>

Quantum Computing

Today's computers, both in theory (Turing machines) and practice (personal computers) are based on classical bits which can be either 0 or 1 to perform operations. Modern Quantum Computing systems operate differently as they make use of quantum bits (qubits) which can be in a superposition state and entangled with other qubits [1]. Superposition and entanglement are thus the two main phenomena that one tries to exploit in quantum computing. Superposition implies that a qubit is both in the ground and the excited state. Entanglement means that two (or more) qubits can be combined with each other such that their states have become inseparable. This gives rise to very interesting properties that can be exploited algorithmically.

The computational power of a quantum computer is directly related to these phenomena and the number of qubits. Two qubits can hold four values at any given time, namely (00, 01, 10, and 11). With each qubit that is added, the compute capacity of the quantum computer is doubled and thus increases exponentially. All these qubits states (in superposition and entangled with each other) can then be manipulated in parallel as, e.g., gates are applied on them which gives the exponential computing power. The problem is that building a qubit is an extremely difficult task as the quantum state that is needed is very fragile and decoheres (losing the state information due to dynamic coupling with the external environment) rapidly. In addition, it is impossible to read out the state of a qubit, which ultimately is necessary to get the answer of a computation, without destroying the superposition state, thus destroying information contained in the qubit state. Basically, it turns into a classical bit that houses only a single value [2].

Current state: A well-known but highly debated example of a quantum computer is the D-Wave machine built by the Canadian company with the same name [2]. It is not yet proven that D-Wave actually uses the above mentioned quantum phenomena nor has any exponential speedup been shown except in one isolated case but which was not considered conclusive by the independent researchers such as M. Schroyer from ETH Zurich [3]. In addition, D-Wave is based on quantum annealing and thus only usable for specific optimization problems.

An alternative direction is to build a universal quantum computer based on quantum gates, such as Hadamard, rotation gates and CNOT. Google, IBM and Intel have all initiated research projects in this domain and currently superconducting qubits seem to be the most promising direction [4] [5] [6] [8].

Currently, the European Commission is preparing the ground for the launch in 2018 of a €1 billion flagship initiative on quantum technologies [9].

Perspective: Making use of Quantum Computing has the benefit to improve the speed-up of certain computations enormously, and even allows solving problems that are impossible for classical computing. Even though the challenges are substantial, they can be separated in physics oriented and engineering oriented ones. The physics challenges primarily have to address the lifetime of qubits and the fidelity of qubit gate operations. The engineering challenges go from identifying relevant algorithms and provide compiler and runtime support. It is also clear that a quantum computer will require a supercomputer to provide the necessary quantum error correction

mechanisms as error rates of around 10^{-3} are not uncommon. As the quantum phenomena require mK (milli Kelvin) conditions, the control logic should be brought as close as possible to reduce the transfer of data up to room temperature computers. Understanding how conventional CMOS behaves under cryo-conditions is another challenge.

Quantum Computing might have the advantage to solve some problems that couldn't be solved with classical computers - one example is Shor's algorithm for decryption which, at least assuming that a large scale quantum computer can be built consisting of millions of qubits, could decrypt a 2000 bit word in around one day which is completely impossible for conventional supercomputers.

In the short term, the Quantum Key Distribution algorithm (QKD) [6] can be used as a new encryption technology that relies on the fact that, when a third party tries to eavesdrop, the entangled state is immediately destroyed.

Further quantum algorithms are [7]:

- Grover's Algorithm is the second most famous result in quantum computing. Often referred to as "quantum search," Grover's algorithm actually inverts an arbitrary function by searching n input combinations for an output value in \sqrt{n} time.
- Binary Welded Tree is the graph formed by joining two perfect binary trees at the leaves. Given an entry node and an exit node, The Binary Welded Tree Algorithm uses a quantum random walk to find a path between the two. The quantum random walk finds the exit node exponentially faster than a classical random walk.
- Boolean Formula Algorithm can determine a winner in a two player game by performing a quantum random walk on a NAND tree.
- Ground State Estimation algorithm determines the ground state energy of a molecule given a ground state wave function. This is accomplished using quantum phase estimation.
- Linear Systems algorithm makes use of the quantum Fourier Transform to solve systems of linear equations.
- Shortest Vector problem is an NP-Hard problem that lies at the heart of some lattice-based cryptosystems. The Shortest Vector Algorithm makes use of the quantum Fourier Transform to solve this problem.
- Class Number Computes the class number of a real quadratic number field in polynomial time. This problem is related to elliptic-curve cryptography, which is an important alternative to the product-of-two-primes approach currently used in public-key cryptography.
- It is expected that machine learning will be transformed into quantum learning - the prodigious power of qubits will narrow the gap between machine learning and biological learning [3].

In general, the focus is now on developing algorithms requiring a low number of qubits (a few hundred) as that seems to be the most likely reachable goal in the 10-15 year time frame.

Impact on hardware: An interesting point to investigate is a better hardware architecture supporting the power efficiency of quantum better. If this is too complex,

it should be at least possible to provide a hybrid architecture of both systems enabling to run the simplest sequences of an application as usually on classical computers and the complex ones on quantum co-processors. By doing this, the system performance can be improved during runtime [8].

As pointed out earlier, a quantum computer will always be a heterogeneous computing platform where conventional supercomputing facilities will be combined with quantum processing units. How they interact and communicate is clearly a challenging line of research [7].

References

- [1] Quantum Computing – University of Waterloo: <https://uwaterloo.ca/institute-for-quantum-computing/quantum-computing-101>
- [2] Cade Metz: Google’s Quantum Computer Just Got a Big Upgrade; <http://www.wired.com/2015/09/googles-quantum-computer-just-got-a-big-upgrade-1000-qubits/>
- [3] Tom Simonite: Google’s Quantum Dream Machine, MIT Technology Review, December 18, 2015, www.technologyreview.com/s/544421/googles-quantum-dream-machine/
- [4] Tom Simonite: Microsoft’s Quantum Mechanics. MIT Technology Review, October 10, 2014. <https://www.technologyreview.com/s/531606/microsofts-quantum-mechanics/>
- [5] Tom Simonite: IBM Shows Off a Quantum Computing Chip. MIT Technology Review, April 29, 2015 www.technologyreview.com/s/537041/ibm-shows-off-a-quantum-computing-chip/
- [6] A. Odeh, K. Elleithy, M. Alshowkan, E. Abdelfattah: “Quantum Key Distribution by Using Public Key Algorithm(RSA)” London, United Kingdom: Third International Conference on Innovative Computing Technology (INTECH), August 2013.
- [7] Daniel Kudrow, Kenneth Bier, Zhaoxia Deng, Diana Franklin, Yu Tomita, Kenneth R. Brown, and Frederic T. Chong: Quantum Rotations: A Case Study in Static and Dynamic Machine-Code Generation for Quantum Computers. 2013 International Symposium on Computer Architecture (ISCA-2013).
- [8] <http://www.wsj.com/articles/intel-to-invest-50-million-in-quantum-computers-1441307006>
- [9] <https://ec.europa.eu/digital-single-market/en/news/european-commission-will-launch-eu1-billion-quantum-technologies-flagship>

6. Beyond CMOS

Nanotubes

Carbon nanotubes (CNTs) are tubular structures of carbon atoms. These tubes can be single-walled (SWNT) or multi-walled nanotubes (MWNT). Their diameter is in the range of a few nanometers. Their electrical characteristics vary, depending on their molecular structure, between metallic and semiconducting [1].

A CNTFET consists of two metal contacts which are connected via a CNT. These contacts are the drain and source of the transistor. The gate is located next to or around the CNT and separated via a layer of silicon oxide [4].

Current state: In September 2013, Max Shulaker from Stanford University published a computer with digital circuits based on carbon nanotubes. It contains a 1 bit processor, consisting of 178 transistors and runs with a frequency of 1 kHz. [2]

Nanotube-based RAM is a proprietary memory technology for nonvolatile random access memory developed by Nantero (this company also refers to this memory as NRAM) and relies on the effect that nanotubes lying cross over can either be touching each other or are slightly separated, depending on their position. A NRAM "cell" consists of a non-woven fabric matrix of CNTs located between two electrodes. The resistance state of the fabric is high (representing "off" or "0" state) when (most of) the CNTs are not in contact and is low (representing "on" or "1" state) vice versa. To switch the NRAM between states, a small voltage greater than the read voltage is applied between top and bottom electrodes. In theory NRAM can reach the density of DRAM while providing performance similar to SRAM. [5]

Perspective: It will take a an unknown number of years before NRAM drives might be in production stage [3].

Impact on hardware: CNTs can be utilized for a lot of different applications in several areas of research. The most promising ones for HPC are the construction of carbon nanotube field-effect transistors (CNTFETs), nanotube-based RAM (or Nano-RAM) and the improvement of chip cooling. CNTs are very good thermal conductors. Thus, they could significantly improve conducting heat away from CPU chips [6].

References

- [1] https://en.wikipedia.org/wiki/Carbon_nanotube
- [2] Max M. Shulaker (Stanford University, Stanford) et al.: Nature, <http://www.nature.com/nature/journal/v501/n7468/full/nature12502.html>
- [3] <http://www.computerworld.com/article/2929471/emerging-technology/fab-plants-are-now-making-superfast-carbon-nanotube-memory.html>
- [4] Lorraine Rispal: Large Scale Fabrication of Field-Effect Devices based on In Situ Grown Carbon Nanotubes. Dissertation, Technische Universität Darmstadt, <http://tuprints.ulb.tu-darmstadt.de/2021/>
- [5] <https://en.wikipedia.org/wiki/Nano-RAM>
- [6] <http://www.extremetech.com/extreme/175457-this-carbon-nanotube-heat-sink-is-six-times-more-thermally-conductive-could-trigger-a-revolution-in-cpu-clock-speeds>

Graphene

In 2010 two physicists at Manchester University in the U.K. shared a Nobel Prize in Physics for their work on a new wonder material: graphene, a flat sheet of carbon with the thickness of a single atom. Konstantin Novoselov and Andre Geim discovered the material by applying plain old sticky tape to simple graphite [1].

Graphene grows on semiconductor i.e. on the surface of a germanium crystal, which is seen as big step towards manufacturability, see [5, 6].

Current state: In 2010, IBM researchers demonstrated a radio-frequency graphene transistor with a cut-off frequency of 100 Gigahertz. This is the highest achieved frequency so far for any graphene device. In 2014, engineers at IBM Research have built the world's most advanced graphene-based chip, with performance that's 10,000 times better than previous graphene ICs. The key to the breakthrough is a new manufacturing technique that allows the graphene to be deposited on the chip without it being damaged [4].

Graphene Project is an EC Flagship project with considerable research efforts in making graphene useful, however, still focused more on the material science perspective than on its potential usage for future computer technology. Graphene is among the strongest materials known and has attractive potential also outside of computer technology, e.g., as electrodes for solar cells, for use in sensors, as the anode electrode material in lithium batteries and as efficient zero-band-gap semiconductors [2].

Perspective: Graphene is a promising technology in laboratory. Due to the fact that the new graphene manufacturing method is actually compatible with standard silicon CMOS processes, it will probably be possible to realize commercial graphene computer chip in future [4].

Since in its current form graphene is not suitable for transistors, researchers have been working on a way to convert it for this use.

Impact on hardware: Graphene has an excellent capacity for conducting heat and electricity.

References

- [1] Moskvitch, Katia. A Graphene Discoverer Speculates on the Future of Computing. [Online] Scientific American, January 23, 2015. <http://www.scientificamerican.com/article/a-graphene-discoverer-speculates-on-the-future-of-computing/>.
- [2] Rodewald, Mike. Researchers discover method for mass production of nanomaterial graphene. [Online] November 10, 2008. <http://newsroom.ucla.edu/releases/method-for-mass-production-of-70969>.
- [3] IBM. Made in IBM Labs: IBM Scientists Demonstrate World's Fastest Graphene Transistor. [Online] February 5, 2010. <http://www-03.ibm.com/press/us/en/pressrelease/29343.wss>.
- [4] Anthony, Sebastian. IBM builds graphene chip that's 10,000 times faster, using standard CMOS processes. [Online] January 30, 2014. <http://www.extremetech.com/extreme/175727-ibm-builds-graphene-chip-thats-10000-times-faster-using-standard-cmos-processes>



[5] <http://hackaday.com/2015/10/14/graphene-grown-on-semiconductors-big-step-toward-manufacturability/>

[6] <http://www.nature.com/ncomms/2015/150810/ncomms9006/full/ncomms9006.html>

Diamond Transistors

Diamonds can be processed in a way that they act like a semiconductor. Diamond based transistors can be fabricated.

Current state: Researchers at the Tokyo Institute of Technology fabricated a diamond junction field-effect transistors (JFET) with lateral p-n junctions. The device shows excellent physical properties such as a wide band gap of 5.47 eV, a high breakdown field of 10 MV/cm (3–4 times higher than 4H-SiC and GaN), and a high thermal conductivity of 20 W/cm*K (4–10 times higher than 4H-SiC and GaN). It has been found that this diamond transistor works with excellent electrical characteristics, up to 723 K [1].

Perspective: Currently the gate length of the fabricated diamond transistors is in the single-digit micrometer range. Compared with the current 22nm technology with gate lengths of about 25nm [2], a reduction in size is absolutely necessary in order to allow fast working circuits (limitation of the propagation delays).

Producing reasonable diamond wafers for mass production could be possible with the method of [3]. The time for producing diamond wafers is another factor that has to be reduced drastically to compete with other technologies.

Impact on hardware: The high thermal conductivity of diamond, which is several magnitudes higher than that of conventional semiconductor material, allows faster heat dissipation. This could solve the temperature problem of stacked dies. Switching energy of a diamond based semiconductor is expected to be much smaller than silicon and the maximum operating temperature can be much higher. It may "revive" the traditional Moore law.

References

[1] T. Iwasaki et al., "High-Temperature Operation of Diamond Junction Field-Effect Transistors With Lateral p-n Junctions," in *IEEE Electron Device Letters*, vol. 34, no. 9, pp. 1175-1177, Sept. 2013.

[2] https://en.wikipedia.org/wiki/22_nanometer

[3] Aida, H., Kim, S. W., Ikejiri, K., Kawamata, Y., Koyama, K., Kodama, H., & Sawabe, A. (2016). Fabrication of freestanding heteroepitaxial diamond substrate via micropatterns and microneedles. *Applied Physics Express*, 9(3), 035504.