

Computer Systems

Position Paper for InfoSam2020

Kjetil Nørvåg, Kjell Bratbergsengen, Rolv Bræk, Haakon Dybdahl,
Anne C. Elster, Torbjørn Hallgren, Morten Hartmann, Jon Olav Hauglid,
Svein J. Knapskog, and Lasse Natvig
Norwegian University of Science and Technology
7491 Trondheim, Norway

The context of the computer systems workgroup incorporates research and development in all areas necessary for a complete working computer; from the design of computers and operating systems, to storage and management of data:

- Computer design and architecture
- Storage technologies database systems, and data management
- Operating systems, middleware, and distributed systems
- Convergence of information services and telecom services
- Computer graphics and visualization
- Supercomputing

A common feature of all these areas is the importance of thinking of the system as a whole, and the emphasis on both experimental and theoretical work. The working group has looked at future trends in the sub-areas listed above, and described how these changes will change systems as well as affect the use of the systems.

We will now describe trends in selected areas of computer systems. It should be remembered that computer systems is a very broad research area, and due to lack of space we can only cover some of the topics in this area in this position paper. We refer to the references given at the end of the paper for more detailed analysis of the summary provided in this paper. Finally, we will also outline some issues related to education in the future.

1 Computer design

From a technology perspective, today's computer system are based on semiconductor technologies. Development of these technologies have been providing a lower cost-per-function for the last four decades. Device integration, speed and functionality has increased while the power consumption per unit of computation has decreased [20].

Even though the core fundamentals of computer design will probably remain relatively unchanged for the coming 15 years, it is reasonable to assume that design methodology will focus on continuing efficient utilization of the available technologies in order to maintain momentum in the improvements mentioned above. This task is all but trivial.

Near term (-2010) challenges include power dissipation, low-power, power management, modeling and simulation of high-frequency circuits (5-40GHz) and nanometer structures, design-implementation productivity, design reuse, design integration and test as well as yield enhancement. For a long term perspective (2010-2018) the challenges include use of non-classical CMOS devices and beyond-CMOS devices, global wiring scaling issues, noise management and error-tolerant designs as well as modeling and simulation of new materials [20].

Most pressingly critical are perhaps the challenges of heating due to high power density (according to Vice President and Chief Technology Officer at Intel, the power consumption of microprocessors is expected to reach 18 kilowatts by 2008 [14] if power-saving techniques and power-saving programming is

not employed) and memory latency [13] as well as low-power requirements for embedded mobile applications.

Intra-chip communication and synchronization is costly in terms of latency, power and area. Distributed and parallel designs are potential solutions to the problems, but the lack of programming paradigms for and understanding of such systems remains a major issue. Tighter integration of computation and memory is one example [8] of a more distributed low-level technique. The integration of different types of analog design on the same chip as memory and processor will also continue to increase, in particular in mobile devices [6]; systems on a chip provide low cost, low power and fast-turnaround times.

Research on biologically inspired architectures and methods [32, 19] in literature suggest several novel approaches to some of the challenges, such as error tolerance, adaptiveness, automated design, distributed architectures, massive parallelism etc.

Emerging technologies [19] such as nanoscale [26] devices and quantum computing may in some cases require radically different design methods. It is often claimed that fundamental barriers will very soon end the days of classical semiconductors. However, such barriers have been overcome several times in the past;

”Every time someone does a study on when we will hit those (scientific) limits, we don’t ever hit them. You hit other, more important, financial barriers,

— Lange, IBM Microelectronics.

2 Computer architecture

The Future Needs New Computer Architectures. Contemporary leading processor architectures are extremely complex, resulting in huge development cost and a problematic power consumption. This has been caused by at least three factors:

- The well know processor-memory gap (also called the memory wall), which has given deeper and more complex memory hierarchies occupying more and more of the chip area and power consumption.
- The processors have used a broad range of smart, but often complex, solutions to hide this memory latency, and to exploit a maximum of the parallelism available in the instruction stream (instruction-level parallelism (ILP)).
- The tremendous development in semiconductor technology has made room for such complex solutions and large caches on the processor chips.

To be able to continue offering improved performance in the future, beyond what can be given by ILP, multiprocessors and thread-level parallelism (TLP) must be exploited, especially in the area of throughput computing. This will give an increased interest in known and new ways of exploiting parallelism. The advances in technology are working against the memory wall problem, and the situation begs for new solutions. The combination of increased focus on parallelism and new memory solutions makes parallel computer architecture to an important topic in the future.

New Challenges in the HW/SW Interface -Rethink Compilers. Near future uni-and multiprocessors will exploit both ILP and TLP. We need a better understanding of its combination, and the future role of the compiler in this context. EPIC (Explicitly Parallel Instruction Computing) architectures exemplified by the Itanium Processor Family hand over more work to the compiler [22] and have fostered a renewed interest in compiler research [24]. The memory wall motivates for radically new solutions integrating computations with memory (e.g. computational RAM [5]), but also a renewed interest in older techniques such as SW prefetching etc. [1]. This may result in new languages or libraries for parallel programming [27], also implying focus on the HW/SW interface.

Portable and Efficient Parallel Programming will be Needed, but is Still Missing. ILP and TLP are old, conservative and precautious roads to parallelism. Till now, most programming models have followed one of two extremes: Either (a) designed for the programmer making it difficult for the compiler to exploit the HW, or (b) using explicit HW knowledge in language constructs and directives (pragmas) reducing

programmer productivity and wrecking portability [27, 15]. Computer architects and researchers in programming languages and compiler design should cooperate to find new and portable ways to exploit the parallelism inherent in the applications.

Mastering the Complexity by Better Evaluation Methods and Tools. A major bottleneck in computer industry is mastering of the complexity in new designs. Evaluation methods and tools will continue to play an increasing role. Extensive simulations executed on HPC platforms will offer the designer powerful tools for design space exploration. Advanced solutions such as integrated performance monitors and HW performance counters will give new and better possibilities for performance studies focusing on the HW/SW interplay [27].

New Technology Gives New Solutions, but CMOS will Survive to Year 2020. The rapid advances in technology will make well established techniques less important (e.g. virtual memory). Reconfigurable HW makes radically new architectures possible. Quantum Computing is a hot topic that has attained a lot of positive press coverage, but some researchers have, based on fundamental physical arguments, stated that “for general-purpose computing, the classical computer will always perform much better than any possible quantum solution” [4]. CMOS works so well in rugged environments that we do not believe it will be replaced by “bioware” or quantum computing within the next 15 years [27].

Other Challenges to Future Computer Architecture. A major problem in the society of today is that computer systems are undependable and insecure [27, 21]. To improve fault tolerance and security we need more dependable HW but also new HW solutions and better HW/SW interfaces that help the SW designer to design better systems.

3 Operating systems

When considering the future of operating systems, one of the first problems is defining exactly what the concept of an operating system encompasses. This is made increasingly difficult as the border between applications, operating system and hardware becomes ever more unclear. For example, the Palladium security technology planned for the next Windows family (codename Longhorn) relies on specific processor extensions, and traditional third-party tools such as firewalls and antivirus programs will soon become included with most major operating systems. A different, but equally important, concern is the nature of the hardware platform the OS is set to govern. The huge span from small embedded systems via traditional desktop computers to large supercomputers indicates that a discussion about general challenges for future operating systems is difficult. As operating system development is mostly driven by industry rather than research, the following discussion is more geared towards what we think will happen rather than what we would like to have.

To first address the future development of core technologies (process handling, storage management, I/O, etc) for traditional computers, several trends are apparent. First of all, the continued development of external devices requires the operating system to keep up. Key features in this respect is plug-and-play and self configuration. Another trend is improved support for large server architectures (as manifested by the new 2.6 kernel for Linux). As for file systems, the planned Windows Future Storage (WinFS) extension to the current NTFS represents an interesting merge of traditional file systems and database technology.

For commercial desktop operating systems, the user interface plays a vital role. Current developments show that improving usability and supporting new input and output devices will continue to be important areas. Tied to improving usability is the extension of the operating system to include extra services. Security issues (viruses, worms) and SPAM presents serious threats to current and future usability of computers and will have to be taken into account.

Light-weight operating systems can be expected to become much more commonplace within the next decade due to the increased use of non-traditional computers such as embedded systems, PDAs, TabletPCs, etc. These systems have a somewhat different set of requirements compared with traditional computers.

Key areas include wireless support, power saving, resource management, device control, real-time operation and reliability.

For more details we refer to [12, 28, 29, 35].

4 Storage technologies, database systems, and data management

The future and development of data management will depend heavily on available storage technologies. Currently, temporary data is stored in volatile RAM, while persistent data is stored on hard disk. Additionally, in applications with large data volumes some data is stored on tertiary storage like tape, CD, or DVD, and in the case of mobile devices, Flash memory is used. The future of non-volatile storage is hard to predict. For example, the death of tape storage has been announced repeatedly. Still, tape is heavily used, and capacity is increasing (for example, 1TB cartridges demonstrated in 2002, probably in market in a couple of years from now [25]). For many years now, variants of high-capacity random-access storage have been expected to be out on the market "real soon", but the development seems to take time. Examples of such technologies are IBM Millipede [18, 7] and Thin Film [30] which both should provide terabyte storage on a single chip. However, if these technologies make it to the market, the impact will be heavy:

- Everybody can afford to store enormous amounts of data. For example, much of the text now only available on the web can be stored on a small device, available to local search and mining. Large storage capacity also means that we do not have the same need for deleting data as before; we can more often store the complete history of all current and deleted data/files.
- Query execution, which is now quite mature and where now mostly incremental performance improvements are done, will have new opportunities. The most important performance bottleneck now is disk arm movement, and algorithms and strategies have as the primary goal to reduce the impact of this bottleneck. If storage technologies with fast random access make it to the market, other algorithms and strategies can be more appropriate to utilize from the new technologies.
- Data transfer between the storage chips and the rest of the system can become an even larger bottleneck than today; more capacity on both system buses and external networks will be needed. The bandwidth of the new storage chips will also be an issue.
- Actually managing and searching the stored data will be a problem. The problem of searching in large amounts of data has traditionally been solved by indexing. If we expect storage capacity to increase at a much higher rate than storage bandwidth and processing (searching) capacity, new indexing techniques might be necessary. They will probably focus less on space efficiency than current techniques. It is possible that more queries have to be pre-computed and stored together with the data in order to benefit from the storage capacity.

The traditional "core" database field has now reached a high level of maturity. Although database systems will continue to become faster, more reliable and more user-friendly, the real challenge in the coming years will be to make them more suitable for the next generation of applications [3]. More general means a broadening from databases to data management in a broader context. This includes, for example, mobile databases (which often are relatively small), but the real challenge is the management and access of the large amounts of distributed heterogeneous data available in one form or another on the WWW. The importance of interoperability will in this context continue to increase, partly solved by the use of XML as storage and transfer format (more and more business data are now stored in, use in 2004 estimated to 60% [34]) and the use of web services (already available) and the semantic web (if it succeeds).

It is also possible that the architecture of many database (and data management) systems in the future will be very different from current technologies. We will probably see heavy distribution, no fixed home location of data, and access based on authorization and roles. Many systems will be based on peer-to-peer-like architectures (it should be mentioned that such systems already exist in a primitive form: for example, a system like Kazaa already indexes 4600 TB of data in 600 million files, distributed between over 3 million users).

Reliability and availability might be achieved by replicated storage. However the problem of guaranteeing correct data has to be solved by different means. More and more data are stored over a long

period of time. Data are eroding both physically and logically. No storage media holds data for ever. To avoid loss of data, programs for rewriting or transferring data to new media has to be established. Logical erosion is more subtle. Formats might change, however this is taken care of by the rewriting programs. Changing of terminology, interpretation, etc is more difficult. Data should be interpreted according to the time of registration. This means handling aging ontologies and historical data correctly.

5 Convergence of information services and telecom services

Open Platforms for Convergent Services. For the establishment of end-user services in a scenario with converging technologies, there has been shown a growing interest to work with open interfaces and standardized service capabilities which opens the value chain and paves the ground for third party service developers and service providers. Although platforms for information services and telecom services traditionally have been different, there are synergies to be efficiently exploited by joint efforts from developers from the previously separated ideological camps. Convergent platforms will need to support both, which means that in the near future, peer-to-peer technology will be integrated with client-server technology, and security aspects (availability, confidentiality, integrity) will become increasingly important. In the subsequent stages, an even more close integration of end-user service design and development and network service provision will be expected, with increasing emphasis on quantitative Quality-of-Service (QoS), including information security, and negotiated Service Level Agreements (SLA).

Service and application architectures. To serve the open platform service provision scenario efficiently, application and service architectures and execution frameworks need to become far more dynamic than today, allowing services and applications to be incrementally developed and dynamically deployed while maintaining high levels of QoS with emphasis on availability and security. These requirements also call for the timely development of compositional methods and tools for dynamic and quantitative verification and validation principles, methods and tools. One possible future benefit may be the successful development and introduction of self configuring systems (autonomic systems in IBM parlance), reducing the perceived user interface complexity and establishing adequate trust, and thereby lowering the end-user acceptance threshold for new services.

6 Computer graphics and visualization

The science of computer graphics has evolved since the 1960s and a solid scientific and methodical base has now been established. Only recently, however, have sufficiently powerful computers at affordable prices been brought to the marketplace making it possible for researchers to better understand the data produced by their applications.

Convenient and intuitive interaction with virtual reality systems is an important challenge. It will increasingly be a demand for tools of interaction that are small in physical dimensions and that can respond to hand and body gestures, to eye movements, to voice and even maybe to facial expressions. Virtual environments of today are mostly restricted to visual expressions. Spatially correct sound is entering the scene. It is expected that smell and other sensations will be part of the future virtual environment. Especially two market segments will be goals for new development, namely the home market and the mobile platform market.

Many aspects of life will be influenced by the virtual reality technology. Only some examples can be mentioned. A very important application will be in education at all levels both for demonstration purposes and for training. Another aspect of the same nature employing the same methods and techniques is entertainment. The word edutainment is increasingly being accepted as a common denomination of these two applications. Virtual heritage is a similar application.

Virtual presence will make it possible to virtually be in the same room for instance for a meeting even if

the participants are on different continents. Research is being conducted in universities. One may also envision that future television will be stereoscopic and immersive. Maybe the fashion catalogue will contain models that the customer may put on a model of her own body and then animate to get a very realistic impression how it will look.

Solutions of mathematical problems and complicated technical and scientific problems may be visualized in a virtual environment. Surgeons may train in virtual environments and perform remote operations using virtual reality techniques.

Among the basic research topics are methods for modeling and animation, parallelization, fast photogrammetry and model synthesis, methods and tools for interaction as well as the physical equipment for virtual environments.

7 Supercomputing

Over the past decades supercomputing systems has evolved from having a focus on large dedicated vector machines like the Cray series of the early 90's to large parallel distributed memory systems based on COTS (commercial off-the-shelf) parts using cluster technologies. The cluster-based systems are generally large groups of PC or workstation processor boards networked together over a local area network or switch. Currently, the "Top 500 Supercomputers" list [31] is dominated by such machines. This trend is expected to continue for the next several years [2].

In addition, as previously predicted [9, 10], traditional HPC (high-performance computing) will probably no longer drive processor technology. Instead, the consumer market – which focuses on the gaming and entertainment world – will have a even larger impact the processor market. This trend is already apparent with Intel's February 2004 announcement of their new Pentium processor codenamed Prescott, which includes 13 new instructions several related to gaming and video [33, 23]. Supercomputing software and high-end library and compiler research will hence have to find new techniques to utilize these features, rather than new architectural features being driven by traditional supercomputing applications.

Another trend in supercomputing is the move towards GRID computing – the idea of providing computing power the way electricity is delivered: plug in your computer and run your application without regard to where the processing power and data storage may be located [11]. This requires break-through in inter-operability (incl. login, compilers, etc), security, as well as streamlining user policies across national borders. Large companies such as Hewlett-Packard and IBM [17] are currently commercializing grid technologies.

Tighter coupling with visualization will become increasingly more important in order to make sense of the Petabytes of data streams generated in the future. See the visualization section for more details.

Finally cluster and GRID computing as well as high-end graphics are also evolving in the area of MMOGs (Massively Multi-player On-line Games) where thousands, and eventually millions, of players play games in a virtual world. E.g. Sony's EverQuest has currently over 460 000 subscribers that are supported by a server farm each serving up to 2000 players simultaneously. [16]

8 Computer systems education in 2020

Education in studies related to computer systems in general aims at giving a broad introduction to the whole area during the first part of the study, with further specialization later. This is likely to be the case also in the future, with curriculum based on what is actual technology at the moment.

References

- [1] 36th Annual International Symposium on Microarchitecture,
<http://www.microarch.org/micro36/>.
- [2] G. Bell and J. Gray. What's next in high-performance computing? *Communications of the ACM*, 45(2), 2002.
- [3] M. L. Brodie. Panel: Future Directions of Database Research -The VLDB Broadening Strategy, Part 2. In *Proceedings of VLDB*, 2000.
- [4] Classical vs. quantum computers: And the winner is ...
<http://www.newsfactor.com/perl/story/22109.html>.
- [5] Computational RAM. <http://www.eecg.toronto.edu/~dunc/cram/>.
- [6] D. Edenfeld. 2003 technology roadmap for semiconductors. *IEEE Computer*, 37(1), 2004.
- [7] E. Eleftheriou. A nanotechnology-based approach to data storage. In *Proceedings of VLDB*, 2003.
- [8] D. Elliott. Computational-RAM : The case for SIMD computing in memory. In *Workshop on Mixing Logic and DRAM: Chips that Compute and Remember at ISCA '97*, June 1997.
- [9] A. C. Elster. High-performance computing: Past, present and future. In *Proceedings of PARA 2002*, 2002.
- [10] A. C. Elster. The next level of simulations: Extreme computing. In *SIMS 2002*,
http://ntsai.oulu.fi/Tapahtumat/SIMS_CallForPapers/1016a.pdf, 2002.
- [11] I. Foster. The grid: A new infrastructure for 21st century science". *Physics Today*, 55(2), 2002.
- [12] Future of operating systems: simplicity
<http://www.computerworld.com/softwaretopics/os/story/0,10801,77335,00.html>.
- [13] P. P. Gelsinger. Microprocessors for the new millenium: Challenges, oppotunities, and new frontiers. In *Proceedings of the International Solid State Circuits Conferance, ISSCC*, pages 22–25, 2001.
- [14] P. P. Gelsinger. Power play. *Communications of the ACM*, 45(6), 2002.
- [15] M. Goudreau. Portable and efficient parallel computing using the BSP model. *IEEE Trans. on Computers*, 48(7), 1999.
- [16] <http://everquest.station.sony.com/about.jsp>.
- [17] IBM Grid Computing home page <http://www-1.ibm.com/grid/>.
- [18] IBM Millipede, <http://www.zurich.ibm.com/st/storage/millipede.htm>.
- [19] International Technology Roadmap for Semiconductors. Emerging Research Devices, 2003,
<http://public.itrs.net>.
- [20] International Technology Roadmap for Semiconductors. Executive Summary, 2003,
<http://public.itrs.net>.
- [21] Jim Gray Touring Award Lecture. <http://research.microsoft.com/~Gray/talks/>.
- [22] V. Kathail, S. Schlansker, and R. Rau. Compiling for EPIC architectures. *Proc. of The IEEE*, 89(11), 2001.
- [23] Intel's Prescott, Intel's Extreme Edition, and AMD's Athlon 64 3400+,
<http://www.linuxhardware.org/features/04/02/02/1654234.shtml>.
- [24] Open Research Compiler project. <http://ipf-orc.sourceforge.net/>.
- [25] J. P. Scheible. A survey of storage options. *IEEE Computer*, 35(12), 2002.
- [26] Technology roadmap for nanoelectronics (microelectronics advanced research initiative), 1999
<http://www.cordis.lu/esprit/src/melna-rm.htm>.
- [27] The Future of Computer Architecture Research, Panel, ISCA-2003.
<http://www.comparch.ncsu.edu/future/>.
- [28] The inevitable future of our operating systems -part 1
http://www.thetechlounge.com/article.php?directory=os_future.
- [29] The Microsoft future according to Bill Gates,
<http://www.ameinfo.com/news/Detailed/33384.html>.
- [30] Thin Film, <http://www.thinfilm.se/>.
- [31] Top 500 supercomputers <http://www.top500.org>.
- [32] A. M. Tyrell, P. C. Haddow, and J. Tørresen, editors. *Evolvable Systems: From Biology to Hardware. Fifth Int. Conf., ICES 2003*, volume 2606 of *Lecture Notes in Computer Science*. Springer, 2003.
- [33] O. Valmot. Lynraske prosessorer. *Teknisk Ukeblad*, (7), 2004.
- [34] S. J. Vaughan-Nichols. XML raises concerns as it gains prominence. *IEEE Computer*, 36(5), 2003.
- [35] Windows "Longhorn" Activity Center,
<http://www.winsupersite.com/longhorn/>.