Workshop NVRAM Technologies

Paris, Salle Inria
https://goo.gl/maps/JtnFhA1qaSS2

Lundi 29, Mardi 30 Mai 2017

Organisateurs : Henri-Pierre Charles, Kevin Marquet, Fabrice Rastello

Formulaire d'inscription http://goo.gl/WBZnLj

Table des matières

Lundi 29 mai 2017

13:00 Accueil

13:15 Impact of emerging memories on various fields of computer science ; introduction to the workshop

Kevin Marquet Inria

Abstract : New memory technologies impact various fields of computer science but it is hard to evaluate, as specialists of one scientific domain, what are these fields and how important is this impact. In this talk, I will present this problematics and I'll introduce the various themes that we will talk about during the workshop. Also, I'll give some examples of existing works that let me think that emerging memories will lead to new usages in computer science, and not only to minor optimisations of some domains.

Bio : Kevin Marquet is an associate professor at the Computer Science and Information Technology department of INSA-Lyon, and a member of the CITI laboratory since 2010. He holds a PhD (2007) and a MSc (2004) both from Université Lille 1 where he studied automatic dynamic memory management for embedded Java applications. Before joining the CITI lab, he was a post-doctoral researcher at the Verimag lab in Grenoble, France where he worked on embedded systems verification. His research interests lie in the area of compilation and operating system support for embedded systems, with a focus on memory management.

14:00 From Embedded World to High Performance Computing using STT-MRAM

Lionel Torres LIRRM

Abstract : The scaling limits of CMOS have pushed many researchers to explore alternative technologies for beyond CMOS circuits. In addition to the increased device variability and process complexity led by the continuous decreasing size of CMOS transistors, heat dissipation effects limit the density and speed of current systems-on-chip. For beyond CMOS systems, the emerging memory technology STT-MRAM is seen as a promising alternative solution. This talk shows first how STT- MRAM can improve energy efficiency and reliability of future systems (HPC & Embedded Systems). We will present the potential of STT-MRAM to design non-volatile processor with two interesting capabilities for energy-efficient and reliable embedded systems: instant-on/off and rollback. We will also discuss about a hybrid design exploration flow to investigate the overall performance impact of using STT-MRAM into the memory hierarchy of HPC systems.

Bio : Lionel Torres obtained respectively my Master and PhD degree in 1993 and 1996 from the University of Montpellier. From 1996 to 1997 I was in ATMEL company as IP core methodology R&D engineer. From 1997 to 2004 he was assistant professor at the University of Montpellier, Polytech'Montpellier (Microelectronic design) and LIRMM laboratory. Since 2004 he is full Professor and was at the head of the Microelectronic dpt of the LIRMM from 2007 to 2010. I am now deputy head of Polytech'Montpellier (engineering school of Montpellier) in charge of research, industrial and international relationship. He is at the Head of the Labex NUMEV (Laboratory of Excellence on digital hardware solutions, Environmental and Organic Life Modeling). His research interests and skills concern system level architecture, with a specific focus on Non-Volatile Computing based on emerging technologies, especially MRAM. He leads several European, national and industrial projects in this field. He is involved in different major conference as DATE, VLSI, FPL, ISVLSI, DAC and is (co)author of more than 40 journal papers and 150 conference publications and 7 patents

14:45 Maximize energy efficiency in normally-off system using NVRAM

Stéphane Grosand Yeter Akgul Evaderis

Abstract : Energy efficient computing has become the key to enable the portability of new applications onto mobile devices which need to be always smaller and more powerful. As the technology node shrinks, the leakage current increases exponentially in deep submicron CMOS, so that new strategies are required in integrated systems to save power without limiting the processing performances. One of the solutions is to rely on Non-Volatile Memories (NVM) and their integration within complex systems. This presentation will show how to benefit from NVRAM technologies and overcome their limitations into a Normally-Off architecture.

Stéphane Gros, System Architect:

Bio :

- Stéphane has 5 years of experience in digital design. He started his career at Sigma Designs (USA), and then joined the team, at the CEA (France), who later founded eVaderis. Finally, Stéphane was one of the first employee at the creation of the start-up. He oversees the definition of the system architecture, its implementation, verification and benchmark.
- Yeter Akgul, Memory Design Engineer: Yeter holds a PhD from the University of Montpellier (France) on “Power Management based on Dynamic Voltage, Frequency and Body Bias Scaling on System on Chip in FD-SOI technology” supervised by the CEA-LETI in collaboration with the LIRMM. She started memory design, SRAM and MRAM, at Tohoku University (Japan), the leader laboratory on MRAM. Finally, she was hired at eVaderis as a memory and NV-logic designer."

15:30 Pause

16:15 "NVRAM: New Opportunities for Compilers

Erven Rohou Inria

Abstract : Hardware manufacturers are beginning to produce new efficient non-volatile memory (NVRAM) chips, i.e. memory that does not loose its content when power is switched off. This technology is in particular a major enabling factor for future low-power IoT devices, but it also considered in many industrial sectors. Beyond the non-volatility, NVRAMs also exhibit different characteristics compared to classical RAMs, for example limited endurance, of assymetric latencies. We review existing compiler-related work and propose new opportunities for compilers in the presence of NVRAMs.

Bio : Erven Rohou is a Senior Research Scientist. His research interests include portability and performance of applications in the many-core era, optimising compilers for high-performance processors, processor virtualisation, split-compilation, dynamic binary optimisation, and performance monitoring. He received his PhD from the University of Rennes 1 in 1998 and was a post-doctoral fellow at Harvard University in 1999. He spent nine years at STMicroelectronics, leading the Advanced Compilation group. He co-authored more than 30 articles in refereed journals and international conferences.

17:00 The role of NVM in the future of HPC workloads

Patrick Demichel HPE

Abstract : Our society and industry are facing a large number of MegaTrends; the IIoT "Intelligent Internet of Things” and Machine Learning are some of the most promising. We face an insatiable demand for more compute, more storage; at the same time, we observe many signs that we are reaching some fundamental limits of our old technologies and infrastructures. Our labs

demonstrated a decade ago that we have no choice but to implement a radical and holistic transformation to reach the Exascale frontier at a reasonable power envelop. We have many others challenges, like resilience, growing complexity, cost, etc. HPE is conducted a large research program called "The Machine"; now entering in its final development phase with a group of partners grouped in a large consortium named "gen-Z" . We will explore what are the fundamental bricks enabling this historical transformation of our architectures, in particular the central role of the emerging NVM technologies. We will also consider some implications on how we could solve our most challenging problems in a short future, what are the potential benefits, what are the requirements of code refactoring, etc. We will conclude with some early demonstrators of breakthroughs for certain classes of algorithms.

Bio : Working for HPE since 36 years on computer technologies with focus on scientific domains. As Distinguished Technologist for Enterprise Group in EMEA working on HPC, Big Data and IoT domains, helping in the development and integration of all IT innovations for large systems. Collaborate with HPE Laboratories, and our divisions, partners and customers on the development of our programs "The Machine", "Moonshot" and adoption of all emerging technologies like FPGA, ML, etc. . Previous contribution: worked on IA64 development in HP-labs USA

17:45 Fin de la journée

19:00 Diner Train Bleu / Gare de Lyon

Mardi 30 mai 2017

08:45 Accueil

09:00 L-IOT : A flexible Platform for ultra-low power IoT

Ivan Miro-Panades CEA

Abstract : The Internet of Things is expected to comprise billions of connected devices, many of which will be wireless sensor nodes (WSN) communicating through a network. The nodes are spatially distributed and able to measure physical or environmental conditions while transferring data through a wireless link. In this context, L-IoT is an initiative from the Integrated Circuits and System Division at CEA. L-IoT is a highly flexible integrated and mixed-signal platform, providing high flexibility and dynamic adaptivity at architectural, design and technological levels. At architectural level, L-IoT is partitioned into computing and wake-up parts, providing very fast wake-up latencies and also RF, imager, energy or timer wake-up sources. At design level, asynchronous logic, adequate arithmetic and reconfigurable analog blocks are embedded. Finally, FDSOI technology and Dynamic Back Biasing is used to reach very high energy efficiency in a wide range of applicative constraints.

Bio : Ivan MIRO-PANADES received a M.S. in telecommunication engineering from the Technical University of Catalonia (UPC, Barcelona, Spain) in 2002, a M.S. and his Ph.D. in computer science from the University Pierre & Marie Curie (UPMC, Paris, France) in 2004 and 2008 respectively. He joined in 2008 the CEA-Leti where he is currently a research engineer in digital integrated circuits. His current research interests include Internet of Things architectures, energy-efficient system design, multicore SoC design, and Fmax/VDDmin tracking methodologies on advanced CMOS technology nodes.

09:45 Towards a single-level database architecture on byte-addressable non-volatile memory

Ismail Oukid SAP

Abstract : The rise of Big Data emphasized the importance of memory-centric systems with an ever-increasing need for more main memory, pushing DRAM to its scalability limits. Storage Class Memory (SCM), also known as byte-addressable non-volatile memory, or NVRAM, is a new class of memory technologies that combines the low latency and high bandwidth of DRAM with the density, non-volatility, and economic characteristic of traditional storage media (SSDs, HDDs). SCM can be used as an extension of DRAM or as fast storage, however, both use cases do not leverage its full potential. A third option is to use it as a universal memory, that is, as main memory and storage at the same time. However, using SCM as universal memory brings many challenges, such as data consistency and aggravated corruption risks. In this talk, we first examine the different challenges and show how they are handled by state-of-the-art persistent memory programming models. Thereafter, we identify three fundamental building blocks for enabling SCM-based software. Finally, we show how these building blocks can be assembled to build a single-level database architecture

Bio : Ismail Oukid is a fourth (and final) year PhD student at TU Dresden, working in close collaboration with SAP SE and Intel. His PhD work consists in paving the path for next-generation databases on SCM, which led him to design a novel single-level transactional storage engine, called SOFORT. Through his work he gained in-depth insights into the challenges of SCM and investigated data structure design, memory management techniques, transaction concurrency control, and fail-safety testing techniques for SCM. He received an Engineering degree (eq. MSc) from the Grenoble Institute of Technology (Grenoble INP -- Ensimag) in 2013.

10:30 Breaking the Memory Bottleneck in Computing Applications with Emerging Memory Technologies: a System, Design, and Technology Perspective

Michel Harrand CEA

Abstract : Emerging non-volatile memories (eNVM), as STT-MRAM, PCM, CBRAM and OxRAM, have the potential to revolutionize computer architectures. This talk presents the different emerging non-volatile memory technologies, the performances which can be expected from them and where they can be used within the computer memory hierarchy – namely last level cache, main memory, and storage – for which benefits, and the requirements these usages put on the device technology. It then presents the 3 main internal architectures of such memories, their benefits and applications, as well as the requirements for the device technology.

Bio : Michel Harrand started his career in Matra Espace in 1980, where he designed automatic pilot systems for satellites. In 1985, he joined Thomson Semiconductors, which became SGS-Thomson then STMicroelectronics, where he designed numerous integrated circuits in the microprocessor, telecommunication, memory, and mostly image compression fields, and lead a design team before being appointed Director of the embedded DRAM department in 1996. He joined CEA in 2006 to prepare the creation of Kalray, a start-up designing and selling manycore processors, which he co-founded in 2008 and where he served as CTO. He joined back CEA end 2012 and is currently exploring the design and applications of emerging non-volatile memories. He has served in the ISSCC International Technical Committee from 2001 to 2006, and holds more than 40 patents.

11:15 Pause

12:00 Peripheral State Persistence For Transiently Powered Systems

Guillaume Salagnac Inria

Abstract : Recent years have seen the advent of a new class of embedded systems we refer to as Transiently Powered Systems. A TPS is a tiny battery-less platform powered by an external energy source, be it ambient (i.e. Energy Harvesting) or artificial (i.e. Energy Transfer). Despite having a small energy buffer (e.g. a capacitor) a TPS typically experiences frequent and unpredictable power failures. To address this problem, several techniques have been proposed which allow a TPS to execute a long-running software program even though the hardware reboots many times per second. But these techniques don't work if the application involves non-trivial hardware peripherals (e.g. ADC, UART, or RF transciever). In this work, we design a software mechanism which transparently saves and restores the state of hardware peripherals with no involvement from the application programmer.

Bio : Guillaume Salagnac is an associate professor at the Information Technology department of INSA-Lyon, and a member of the CITI laboratory since 2009. He holds a PhD (2008) and a MSc (2004) both from Université Grenoble Alpes (UGA) where he studied automatic dynamic memory management for real-time embedded Java applications. Before joining the CITI lab, he was a post-doctoral researcher at the CSIRO ICT Centre in Brisbane, Australia, where he worked on high-level programming for Wireless Sensor Networks. His research interests lie in the area of programming languages and operating system support for resource-constrained embedded systems, with a focus on memory management.

12:45 Conclusion