HiPEAC INFO 47

COMPILATION ARCHITECTURE

APPEARS QUARTERLY | JULY 2016

Summer 2016: ACACES summer school, Fiuggi

Marc Duranton: A Vision of the future

Leapfrogging the valley of death with TETRACOM

Cyber-physical systems meet supercomputing

contents

Interview with Marc Duranton

Leapfrogging the valley of death in Europe

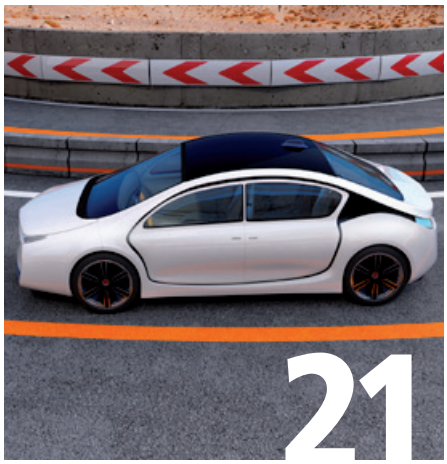

Codeplay: the programming behind your car's eyes on the road

HiPEAC is the European network on high performance and embedded architecture and compilation.

hipeac.net

@hipeac

hipeac.net/linkedin

HiPEAC has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 687698.

Cover photo: Eneko Illarramendi **Design:** www.magelaan.be
Editor: Madeleine Gray **Email:** communication@hipeac.net

Critical systems in avionics: sky-high challenges for multicore development

Technology opinion: The case for open source hardware

Career talk: Mafijul Islam, Volvo AB

Recently, I saw the latest numbers on the cost of ageing in Belgium. The retired population is currently around 20% of the total population and is increasing annually by 0.4%. The total pension cost is already more than 11% of Belgian GDP, and will be close to 12% in 2020. This is the just cost of the pensions, not including the cost of health care, which is on a similar track (now about 7.5% of GDP). Since the Belgian social security system is based on taxes, an increase in social security spending has to be funded by higher taxes, by structural savings in other departments of the government or by economic growth, which happens to be very low at the moment. That means that the growing cost of the retiring baby boomers is affecting all of us in the form of increased taxes and/or reduced benefits. This will continue for the next 15 years when the last baby boomers (including myself) will retire. The budgetary impact can in practice only be mitigated by growing the economy. This situation is not unique to Belgium; many European Union countries are facing similar dilemmas. One way HiPEAC can help is by innovating relentlessly and by helping to digitize European industry. This means that we need to focus more on research challenges with a potential impact on the market.

In April, HiPEAC3 underwent its final review. The reviewers noted a significant and growing positive impact on the computing systems community in Europe. They were also pleased with the professionalization of the communication services of the network. The reviewers encouraged us to keep improving the industrial participation in the network and to continue building links with underrepresented countries in the network. The reviewers gave HiPEAC3 an 'Excellent' final rating. We are very happy with this outcome, and we are committed to continuing our support for the European computing systems community in the future.

This issue of the magazine is released to coincide with ACACES, the annual HiPEAC summer school. The summer school also marks the beginning of the summer break for me and for the HiPEAC staff. We wish you a relaxing summer with your family and friends, and we hope to see you again after the summer holiday in good health, and full of plans for the year to come.

Koen De Bosschere, HiPEAC coordinator

ICT Proposers' Day: what's it all about?

Thinking about submitting a project proposal to a forthcoming European Commission call? Sandro D'Elia, HiPEAC project officer in the Complex Systems and Advanced Computing Unit at the European Commission, explains why the ICT Proposers' Day is an unmissable event.

The ICT Proposers' Day 2016 has just been announced. You can find the programme and practical information on the European Commission website (https://ec.europa.eu/digital-single-market/ict-proposers-day-2016), but here I would like to explain why you should seriously think about going to Bratislava on 26 and 27 September.

The Proposers' Day is a yearly event (which actually lasts not one but two days, just to confuse things) aiming to promote the information and communication technology (ICT) research and innovation programme. It is the place where you can get all the information you need to shape your idea into a real project.

This year the event will be more important than usual, for one main reason: the policy direction which has been given to research and innovation activities. As you may know, Günther H. Oettinger, the Commissioner for Digital Economy and Society, is strongly pushing the initiative for the digitalization of European industry, which will have a direct impact not only on our labs, but on jobs, growth and investment through digital technologies.

This is good news for the HiPEAC community, with its proven expertise in the industrial and professional applications which are the bread and butter of digital industry. The Proposers' Day is the event which explains how the overall policy of the European Commission translates into practical Horizon 2020 actions, both in terms of calls for proposals for 2017 and of general strategy for the coming years.

If you are thinking about submitting a project proposal in 2017, please be aware that European Commission staff are not allowed to discuss proposals in private meetings, so the Proposers' Day is the best opportunity to meet the project officers responsible for the evaluation of proposals, to ask for clarifications about the work-programme, to get feedback on your ideas, and in general to ask any questions you may have about future calls for proposals.

There is also another reason to be in Bratislava, particularly if you are interested in next year's computing call. When the 2016-2017 workprogramme was drafted, due to the limited funding available, the decision was taken to limit the objective of

the 2017 call to software, as you will see from the workprogramme text which is currently published ('Programming environments and toolboxes for low energy and highly parallel computing'). However, due to the stronger policy focus on the technologies for the 'Digitising Industry' programme, and in the context of the European Cloud Initiative, the possibility of allocating extra funding for research on processor hardware is under discussion. A formal decision will be taken in the coming weeks, but at this stage it seems likely that an updated version of the work-programme text will be published in the coming months. Naturally, Bratislava will be the place where you will get all the information about the latest updates.

Join me and my colleagues in Bratislava on 26 and 27 September – we look forward to seeing you there.

"The Proposers´ Day explains how overall EC policy translates into Horizon 2020 actions"

Photo credits: Eneko Illarramendi/Ali Azarian

A Computing Systems Week with Portuguese flavour

Taking place in Porto, Portugal on 20-22 April, the latest HiPEAC Computing Systems Week (CSW) once again provided an opportunity to catch up on the latest developments in high-performance and embedded architecture and compilation.

Hosted by the University of Porto, the event included a total of 11 thematic sessions, which provided an opportunity for lively discussion of key issues in computing systems. Themes included:

- **Ideas for the next HiPEAC Vision**, setting out the future of computing systems. *See our interview with Marc Duranton on p.8.*
- **Models and tools towards energy-efficient exascale computing systems**, presenting European-funded research in different areas. These included message passing, programming, self-adaptivity, reconfigurable architectures, runtime-aware architecture and benchmarking energy consumption.
- **Challenges in the area of teaching parallel computing**, which explored the limitations of current models and tools for teaching parallelism.
- **Technological challenges of the maritime internet of everything**, with presentations focusing on disruptive technologies and their importance for marine sustainability and economic activities.
- **Open source in high-performance computing and curricula**, focusing on hardware, application development and benchmarking. This session provoked heated debate, with the benefits relating to collaboration and exploitation being pitted against potential costs and lack of industry support. *See p.24 for an opinion piece supporting open hardware.*
- **Computing systems industry activities in Portugal**, highlighting some of the variety and ingenuity in Portuguese industry. Companies Coreworks, Altice Labs, Ubiwhere, PLUX Wireless Biosignals, CRITICAL Software and EVOLEO Technologies presented work on technical issues such as signal processing, dynamic bandwidth allocation and mixed criticality, as well as city monitoring, biosensors and spacecraft applications.

Participants learned about the aims and results of a wide range of European Commission-funded research projects. Local organization was provided by João Cardoso and his team at the University of Porto, while Manolis Katavenis (FORTH) was responsible for leading the CSW task. Xavier Martorell and Marisa Gil (Barcelona Supercomputing Center) arranged the thematic sessions. HiPEAC would like to thank the organizers for coordinating an excellent event, and for the warm welcome in beautiful Porto.

ARM acquires Apical, while Intel buys Yogitech

In May, ARM, the low-power chip designer and HiPEAC partner, acquired imaging and embedded computer vision specialists Apical for $350 million. According to ARM, Apical's advanced imaging products are used in more than 1.5 billion smartphones and approximately 300 million other consumer and industrial devices, including cameras and tablets. The move is intended to accelerate ARM's growth into new markets such as connected vehicles, robotics, smart cities, security systems, industrial/retail applications and internet-of-things devices, which require intelligent image processing. Read more on the ARM website: http://bit.ly/ARM-Apical

Meanwhile, in April, chip developer Intel bought Yogitech, represented in HiPEAC by its chief executive Riccardo Mariani. Yogitech, which specializes in semiconductor functional safety and related standards, has joined the Intel Internet of Things Group. Intel stated that this would help them to advance in advanced driver assistance systems (ADAS), robotics and autonomous machines – technology domains which are crucial for certain market areas, such as the automotive sector. In a blog post, Ken Caviasca, the vice president and general manager of platform engineering and development in Intel's Internet of Things Group, explained that the 'industry is now moving from automating data to inform better decisions, to automating actions informed by real-time data', with functional safety being a requirement. The full blog post can be read on the Intel website: http://bit.ly/Intel_Yogitech

Accurate power modelling of embedded processors

Researchers from Electronics and Computer Science at the University of Southampton, in conjunction with ARM, have developed a free open source software tool, PowMon, that provides accurate power models of embedded processors in mobile devices.

The tool was developed as part of the PRiME research programme (www.prime-project.org), a five-year collaborative research programme funded by the UK Engineering and Physical Sciences Research Council, which involved Imperial College London and Manchester, Newcastle and Southampton Universities, led by Professor Bashir Al-Hashimi. Five companies are also participating in the project as industrial partners.
To find out our results, download the paper from the ACM Digital Library: http://dl.acm.org/citation.cfm?id=2901322

Accurately estimating central processing unit (CPU) power consumption is a key requirement for controlling CPUs at run time – to implement energy saving techniques, for example – and also at the design stage when exploring processor/system-on-chip architectures.

According to Professor Al-Hashimi, 'obtaining accurate data from mobile devices can be challenging and more time consuming than using a simulator or desktop/server devices. For this reason, we have developed the PowMon software tool which allows workloads to be automatically run on a mobile device'.

Metrics such as Performance Monitoring Counters (PMCs), temperature, CPU utilization, CPU power and CPU voltage are collected and analysed to produce accurate power models. PowMon's two main features are that it allows users to automate the model-building methodology for their hardware of choice and that it provides power models for specific CPUs, which can be used as reference models for this hardware.

Further information and the free PowMon tool software download is available at: www.powmon.ecs.soton.ac.uk/powermodeling

Thermal covert channels: red-hot security issues

Davide Basilio Bartolini

Covert channels, which allow ways for information to be communicated not allowed by a computer's security policy, are a well-known security issue. Today, many multicore processors feature easily accessible temperature sensors providing information for dynamic thermal management, which can be exploited to establish a thermal channel: a clear security threat for mobile devices.

In a paper published at EuroSys 2016, we investigate the capacity of thermal covert channels. We use a new methodology leveraging theoretical results from information theory and experimental data to study thermal covert channels on multicores. Spectral techniques are used to analyse data and estimate the capacity of channels from a source application to temperature sensors.

To find out our results, download the paper from the ACM Digital Library: http://dl.acm.org/citation.cfm?id=2901322

Mateo Valero awarded Creu de Sant Jordi

On 26 April, HiPEAC co-founder Mateo Valero was awarded the Creu de Sant Jordi (St George's Cross) by the Catalan Government 'for his work in the field of computer architecture, for which he has received prestigious international awards including the Eckert-Mauchly Award in 2007'.

The Creu de Sant Jordi is one of the highest distinctions awarded by the Catalan Government. It was established in 1981 with the aim of recognising those who have used their talents to provide outstanding services to Catalonia.

Originally from Aragaon, Valero, who has lived in Catalonia since he was 21, commented: '[Catalonia] is where I've spent most of my academic and professional life, where I started a family and put down roots, and where I've been able to forge ahead with many projects. Being awarded the Creu de Sant Jordi is a great honour for me.'
The Creu de Sant Jordi follows the award of the Spanish Association of Electronics, Information Technology and Digital Content (AMETIC) gold medal to Valero during the 2016 Mobile World Congress. In May, Valero was also given an honorary doctorate by the University of Granada.

Making Europe smarter

How can European industry become truly digital? Through a series of innovation initiatives, the European Union-funded Smart Anything Everywhere (SAE) programme supports both high- and low-tech sectors in digital value creation. A workshop held on 13 June, organized by HiPEAC in collaboration with the European Commission, showcased results, best-practice examples and lessons learned from the programme's innovation hubs, before highlighting themes in the Horizon 2020 ICT-04-2017 SAE call, which closes on 8 November.

This was followed by a collaboration workshop on advanced computing and cyber-physical systems (CPS) in which over 40 projects were presented, covering technical themes including next-generation servers, modelling, toolchains and methods for CPS, heterogeneous computing, real time/reliability, cloud and mixed criticality, as well as community building and roadmapping. The Commission showed how it supports intellectual property rights (IPR) management through the IPR helpdesk (www.iprhelpdesk.eu) and how the Innovation Radar is supporting innovation in Europe (http://bit.ly/Innovation_Radar).

Dates for your diary

NiPS Summer School 2016: 'ICT-Energy: Energy consumption in future ICT devices'
13-16 August 2016, Aalborg, Denmark
www.nipslab.org/summerschool2016

ICT-Energy Science Conference 2016: energy efficiency and sustainability in ICT
16-19 August 2016, Aalborg, Denmark
www.ict-energy.eu/scienceconference_home

25th International Conference on Parallel Architectures and Compilation Techniques (PACT)
11-15 September 2016, Haifa, Israel
http://pactconf.org/
Including SeedHPC - Seeding Future HPC Programming workshop, endorsed by HiPEAC
https://sites.google.com/site/seedhpc

ARM Research Summit
15-16 September 2016, Cambridge, UK
www.arm.com/summit

24th International Conference on Real-Time Networks and Systems (RTNS)
19-21 October 2016, Brest, France
Call for papers: deadline 25 July 2016
http://rtns16.univ-brest.fr

The 13th IFIP International Conference on Network and Parallel Computing (NPC 2016)
28-29 October 2016, Xi'an, China
www.cvnis.net/npc2016

IEEE Nordic Circuits and Systems Conference (NorCAS 2016)
1-2 November 2016, Copenhagen, Denmark
Submission deadline: 15 August 2016
www.norcas.org

9th IEEE/ACM International Conference on Utility and Cloud Computing (UCC 2016)/ 3rd IEEE/ACM International Conference on Big Data Science, Engineering and Applications (BDSEA 2016)
6-9 December 2016, Shanghai, China
http://computing.derby.ac.uk/ucc2016/

13th International Conference on High Performance and Embedded Architectures and Compilers (HiPEAC 2017)
23-25 January 2017, Stockholm, Sweden
www.hipeac.net/2017/stockholm

"We have to decide how technology will change our lives"

In this interview, Marc Duranton, senior member of the Architecture, IC Design and Embedded Software Division department of CEA (the French Alternative Energies and Atomic Energy Commission), and leader of the HiPEAC Vision task, reflects on the past and future of computing systems. Prior to joining CEA, Marc spent over 23 years at Philips, Philips Semiconductors and NXP Semiconductors. He has MSc degrees in electrical engineering and in computer science as well as a PhD in signal and image processing.

When I was a kid, I was fascinated by cybernetics: the robot 'tortoises' created by William Grey Walter, the self-adapting Homeostat designed by W Ross Ashby, the work of McCulloch and Pitts on artificial neurons and the work of cybernetics pioneer Norbert Wiener. All of these I discovered through books in a library in a provincial French town. Naturally, I became interested in the first microprocessors and computers. I was one of the winners of a French national contest on computing science in 1979 before being awarded degrees in electrical engineering and computer science. My interests ever since have lain both in hardware and in software, as well as in the synergy between them.

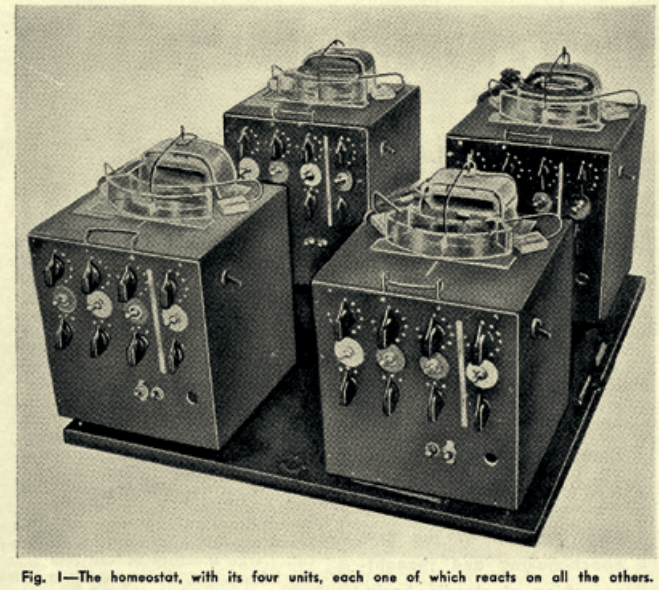

Fig. 1—The homeostat, with its four units, each one of which reacts on all the others.

Ashby's Homeostat. *Reproduced with permission of the Estate of W. Ross Ashby*

Early in my career I joined Philips, first working on 3D graphics accelerators, then on neural networks. I led the design of the L-Neuro chip family architecture – digital processors using neural network techniques. After that, I moved to California, where I worked on video coprocessors for the very long instruction word (VLIW) processor TriMedia and for various Nexperia platforms, including the first Philips system-on-chip for the Japanese digital high-definition (HD) standard BS-Digital. I then came back to Europe to join NXP Eindhoven, where I led the Ne-XVP project on the design of a multicore processor for real-time applications and consumer video processing, before returning to France to work in the CEA Research and Technology Department.

I am currently interested in deep neural networks (now back in business!), in the internet of things and in high-performance computing, as well as leading the HiPEAC Vision activity.

Among my career highlights are the time when L-Neuro chips were sent into space in a satellite, and when we got the first full HD images from our BS-Digital chip. I also have good memories from working with Japanese companies on this chip.

Humans have to decide whether technological developments will change our lives for the better or for the worse. The most exciting evolution of computers is in how they interface with the real world and with people, as well as the increasing ability of machines to 'understand' the environment and the context – not just numbers. Voice recognition has been integrated into all smartphones and is appearing on home devices, while smart cameras have started to interpret what is going on. New algorithms are making machines smarter, taking applications

"Computers can increasingly 'understand' the context – not just numbers"

such as self-driving cars out of the realms of science fiction and making them a reality. Deep learning and cognitive computing will change how we use computers and will open up a whole new range of applications. We will interact with machines in a natural manner, and they will seem to us more and more like 'beings' than 'things'.

Bilateral exchanges between the HiPEAC community and the European Commission are very important: through the HiPEAC Vision, we can inform the European Union about forthcoming evolutions and their potential societal impact – positive or negative – and help them define research priorities. For their part, they can set work programmes and open interesting calls for projects.

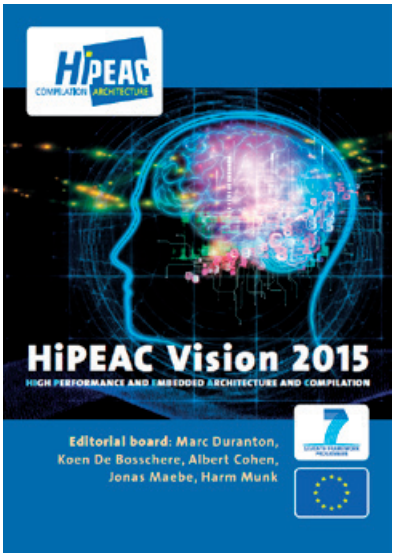

The last HiPEAC Vision was published in 2015

We need a new HiPEAC Vision for 2017, because technology is evolving rapidly. There are new developments constantly: new cognitive computing systems, for example, such as Google's AlphaGo, which managed to beat champion Lee Se-dol at the Chinese board game Go, as well as new applications and ever-changing security issues. New technologies like quantum computing are coming into focus and other are hitting their limits. There is also the need to provide technologies to tackle the societal challenges we're facing.

Among the challenges ahead are that **computing systems will need to interact with the physical world** to understand what is going on, to take decisions, and to help us in our everyday lives. To allow this to happen, we need to trust the systems, with all the issues that entails: we need systems which are reliable, which will not crash, which are safe from attacks by hackers. We also need systems which are highly energy efficient, both for embedded computing in internet devices or sensors which will need to have an extremely long battery life or scavenge energy from the environment, and for supercomputers and servers, which will need to consume less energy and dissipate less power so they can be affordable to run.

The broader the range of contributions to the HiPEAC Vision, the better it will be. Of course we need input from the HiPEAC community, most of whom are computing systems researchers, but also from application developers and people who use computing devices without knowing how they work. We're entering the era of 'vanishing computers', where computers are no longer represented by a screen and a keyboard: they are everywhere and they are invisible. They will need to be made transparent for users so that these devices can enhance quality of life. We will talk to the computers, they will see what is going on; this might be the next step after the era of the keyboard/screen and the touchscreen era. All contributions will be important in helping shape our future.

"We're entering the phase of 'vanishing computers'"

The HiPEAC Vision has evolved since its first incarnation in 2008. **In the beginning, it was a roadmap rather than a 'vision'**: a highly technical document aimed at specialists setting out steps for future developments. Over time we saw the need for a broader scope which would take not only technology but also the impact of computer science on society into account. We need this wider focus to really tune into what the technology will be. Things can happen very quickly: it's worth remembering that just 10 years ago there were no smartphones and now they're everywhere and have changed how we live. Things are evolving especially rapidly from the perspective of user acceptance: the technologies to create smartphones have existed for some time, but it was the integration of these technologies onto a single device, the business model and the adoption of this device by users which caused them to become widespread. Hence the HiPEAC Vision is really a combination of what's happening on the technology side, including current limitations, along with application drivers and societal requirements.

Download the Vision for free from the HiPEAC website www.hipeac.net/roadmap

Contribute to the next HiPEAC Vision – respond to the survey: www.surveymonkey.com/r/hipeacvision2025

Leapfrogging the vall

You're a research institute with a great idea but no means of making sure it reaches the market, or you're a small enterprise with plenty of business expertise but no budget for research and development. Where do you turn? European Union-funded programmes such as TETRACOM are here to help, as these inspirational examples show.

TETRACOM: A SPRINGBOARD FOR START-UPS

Rainer Leupers, RWTH Aachen University, TETRACOM Project Coordinator

The TETRACOM coordination action (www.tetracom.eu) provides new incentives for academia-industry technology transfer via focused, bilateral technology transfer projects (TTPs) across the whole of Europe. A total of 50 TTPs are being supported and co-funded, covering a multitude of ICT areas, such as the automotive sector, communications and multimedia, data analytics, health, industry 4.0, as well as safety and security. Each TTP brings together one academic partner and one industry partner and implements the transfer of a particular hardware or software technology or intellectual property (IP). The vast majority of TTPs involve industry partners from small/medium enterprises (SMEs), which have reported major benefits from TETRACOM through new products, significant cost savings, or improved processes.

Being a co-founder of several companies myself, I'm particularly glad that many TTPs have also helped European start-ups get off the ground by transferring key technologies that contribute to the core of their product offer. This article provides highlights of three such start-up companies, all of which have been supported by TETRACOM: Bluebee, a provider of highly efficient genomics analysis cloud services; Creonic, which offers advanced hardware IP blocks for 4.5G communications; and Xsensio, which is developing a groundbreaking 'lab-on-skin' device for next-generation healthcare. I'm excited to see this level of industrial impact from a European project and I cordially wish them long-term market success.

"SMEs report benefits from TETRACOM through new products, cost savings, or improved processes"

ley of death in Europe

ACCELERATING GENOME-BASED DIAGNOSTICS FOR CLINICAL USE

Zaid Al-Ars and Koen Bertels, Delft University of Technology (TUDelft)

Delft-based Bluebee was founded by (left to right) Vlad-Mihai Sima, Koen Bertels and Zaid Al-Ars

The rapid increase in throughput of DNA sequencing technology, coupled with the exponentially decreasing prices of DNA data acquisition, have resulted in a computational bottleneck in the analysis phase of genomics analysis pipelines. This bottleneck is hindering the deployment of these analysis techniques for clinical use to diagnose genetic diseases and propose personalized therapies based on the DNA profile of an individual.

To overcome this, we implemented a scalable, high-performance computational infrastructure specifically targeted towards the field of genome analysis, which enables hospitals, clinics and research institutes to manage their computational challenges. We augmented this high-performance infrastructure with field-programmable gate array (FPGA) and graphics processing unit (GPU)-accelerated pipeline components to improve the cost efficiency of the process and increase the throughput of the analysis.

These hardware-accelerated implementations of computationally intensive kernels run up to an order of magnitude faster than regular processors. TETRACOM funds were used to transfer some of the intellectual property developed at TUDelft for these kernels to our spin-off, Bluebee, where the kernels were integrated into the total application pipeline.

Commercial potential

The field of genomics is extremely promising, with huge potential to influence our society and revolutionize the medical domain. However, until recently it has been mostly restricted to the scientific community, addressing biological research questions. Over the last couple of years, a number of scientists have started using their research in clinical applications to diagnose and treat genetic diseases such as cancer. Deploying such computationally intensive techniques in a production environment sets stringent performance and cost requirements and opens up major potential for commercialization.

Certainly, the scalable, accelerated solutions which we've developed show the viability of these computational approaches to high-performance genomics. This is thanks to the abundance of potential parallelism that genomics applications contain, and the high computational intensity of the algorithms. Various genomics-related applications would benefit from the potential performance improvement achieved by applying scalable and accelerated computational techniques.

> *"Our scalable, accelerated solutions show the viability of computationally intensive approaches to high-performance genetics"*

In our case, we found that the results of the integrated EU-funded project hArtes (Holistic Approach to Reconfigurable Real Time Embedded Systems), of which TUDelft was the scientific coordinator, were too good to remain unused. Hence we established a spin-off, Bluebee, which offers services based on technology directly related to the project. In the first phase, the emphasis of Bluebee was on trying to sell the core technology. However, as that turned out to be very difficult, we then made the switch to developing an end product. Although we initially undertook a very successful pilot in the financial industry, we

decided that computational genomics had more potential and was starting to take off.

The genomics analysis platform being developed by Bluebee provides a scalable, accelerated and secure private-cloud solution targeted at clinical application of genomics analysis techniques. The platform's scalability allows seamless increase of the throughput (i.e., allows more patients per hour to be analysed) when required by the clinic. Acceleration, on the other hand, allows for the latency of the analysis to be reduced (i.e., allows faster sample-to-diagnosis for a specific patient). Meanwhile, the platform's security features make it suitable for clinical use where privacy needs to be assured by law.

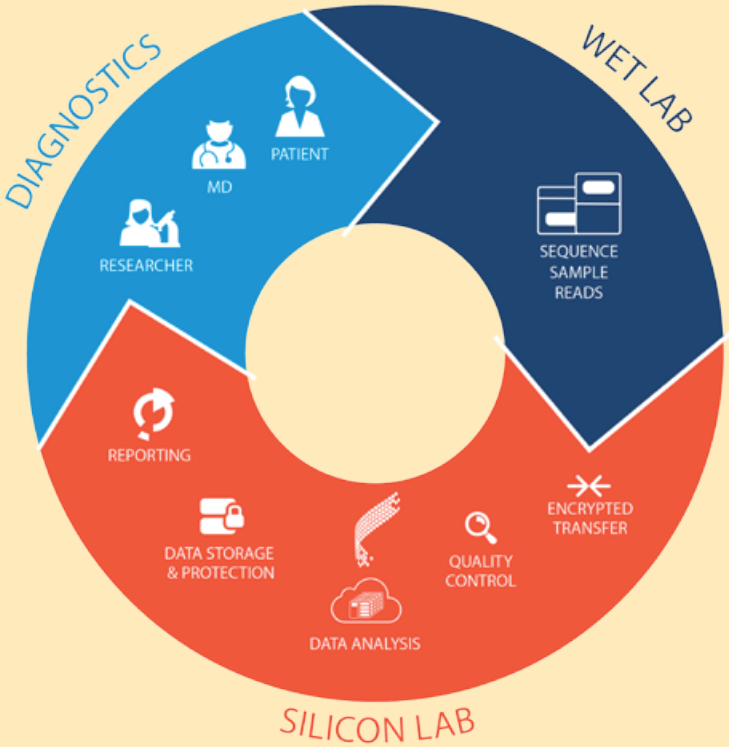

Bluebee private cloud

Launching a successful start-up

While relatively young, the market for clinical genomics has immense growth potential and is expected to have a significant impact on various aspects of daily life. Early players in the field have a strong influence on market development and are well positioned to capture a sizeable share of this growing market. In addition, the field of genomics focuses on creating value for society, specifically in the medical domain, an effort that is appreciated by all. 'Good science', as the slogan of Bluebee indicates, represents the ethos as well as the activities of the company.

Important factors when convincing investors to back Blubee were the successful proof of concept and the fact that we managed to attract a highly experienced management team who were enthusiastic about the concept from the outset. It is no exaggeration to say that those were the game changers for Bluebee.

> *"A successful proof of concept and highly experienced, enthusiastic management team were game changers for Bluebee"*

Some of the most important advice we could give to young entrepreneurs would be to follow your passion, to focus your efforts, and to keep an eye out for commercial opportunity. Fortunately, there is an increasing number of European funding programmes, such as TETRACOM, that can help ambitious scientists capitalize on their intellectual worth. Hopefully, we will have more programmes in this area to encourage more entrepreneurial activity in Europe and enhance Europe's global competitiveness.

A TURBO CODE DECODER FOR NEXT-GENERATION MOBILE DEVICES

Christian Weis, University of Kaislerslautern

In addition to other research activities, the Microelectronic Systems Research Group at the University of Kaiserslautern has a great deal of experience (totalling more than 250 'person years') in designing and verifying high throughput channel decoders. This work centres especially on two high-performance, forward-error correction codes for efficient information transfer over communication channels: turbo codes and low-density parity-check (LDPC) codes. Essentially, these two types of code allow a high proportion of data to be transmitted reliably in the presence of data-corrupting noise.

Founded in 2010, Creonic is a spinoff of the Microelectronics System Research Group which designs and delivers intellectual property (IP) cores for integration in application-specific integrated circuits (ASICs) and field-programmable gate arrays (FPGAs). The focus is on complex signal processing functions for communication systems ranging from digital video broadcasting to LTE-A, the 'advanced' enhancement in the Long-Term Evolution (LTE) wireless communication standard.

The Microelectronic Systems Research Group at the University of Kaiserslautern (left) worked with the Creonic team (right) to deliver new LTE turbo IP

In order to be competitive, Creonic needs to be able to provide the best LTE turbo IP core in the area of very largescale integration (VLSI) and communications performance. Through a technology transfer project supported by TETRACOM, the University of Kaiserslautern worked with Creonic to adapt the LTE turbo code decoder IP to the needs of Creonic's customers. The target markets centre on mobile chipsets and base stations – the transceivers which connect devices to one another – for 4.5G, the latest operational generation of mobile communications.

The main improvements with regard to the state of the art are smaller chips, which result in lower energy consumption and longer battery life. The first integrations for early customers have now been implemented, and further developments based on the current design will be made this year. The design will provide the basis for further developments in the context of the European Union's Horizon 2020 research and innovation programme and 5G mobile telecommunications.

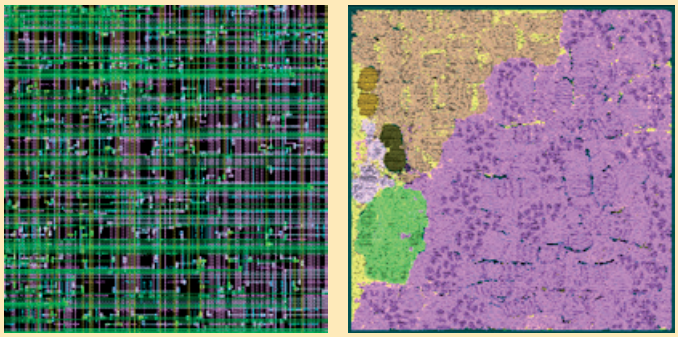

Layout plot details (left) and hierarchy (right)

Conditions within the European Union are not currently favourable for technology transfer, as many small and medium enterprises (SMEs) simply cannot afford to incorporate additional research and development activities to foster innovation. Providing additional support is therefore essential to increase the momentum of technology transfer. By allowing cost-efficient transfer of knowledge, TETRACOM lowers the commercial risk of developing innovative products in young companies. Without the support of TETRACOM, it would have been not affordable for Creonic to develop the LTE IP in such a short timeframe.

"Without TETRACOM's support, it would have been not affordable for Creonic to develop the LTE IP in such a short timeframe"

ANALYSE YOUR SWEAT WITH XSENSIO'S LAB-ON-SKIN

Esmeralda Megally, Xsenio, and Adrian Ionescu, École polytecnique fédérale de Lausanne (EPFL)

The Nanoelectronic Devices Laboratory (NANOLAB) is a lab at EPFL in Switzerland that focuses on developing, designing and modelling energy-efficient nanoscale solid-state devices. NANOLAB has been at the forefront of European efforts to develop new sensing devices at the nanoscale. In 2014, Xsensio was founded as a spinoff of NANOLAB to commercially exploit novel field-effect transistor (FET) sensors in order to develop a wearable tool which senses biochemical information on the skin's surface (from electrolytes to metabolites, small molecules and proteins) to provide unprecedented real-time information about our health and wellness, in a simple and non-invasive way.

Sweat on the skin's surface can provide a wealth of information that is currently not exploited to support health care and wellness applications. Until recently, there was no technology available to continuously collect sweat and analyse its composition in real time, but recent advances in miniaturized sensing and computing have paved the way for a fundamental paradigm shift. The Lab-on-Skin™ chip being developed by Xsensio will collect infinitesimally small volumes of sweat and conduct repeated multi-parametric analysis in real time to provide meaningful information to the user.

Xsensio was looking for a novel on-body interface that would continuously collect and analyse sweat based on a multi-parameter sensing system. The challenge was to develop a technological solution combining nanofluidics and multiple sensing capability in a single system on chip with ultra-low power consumption. In particular, the smart chip should be able to operate in a wearable form 24 hours a day, seven days a week for many days in a row, which implies significant progress beyond the existing industrial state of the art.

Passive channeling of a drop of liquid through pumpless nanofluidic channels

Xsensio approached NANOLAB with a view to working in partnership on this project. Following multiple iterations, a common win-win path for joint development and innovation was identified.

Groundbreaking health applications

NANOLAB is developing a highly innovative zero-power, nanofluidic-on-sensor interface, which allows infinitesimally small sweat droplets to be pumped and channelled through its many capillary channels. The solution, which is based on organic and biocompatible materials that can be processed on silicon wafer carriers, neither relies on external power sources nor uses a pump.

Xsensio will now integrate this unique nanofluidics interface into a wearable form and adapt its functionality so that it can be used to analyse a range of different chemicals. The final result will be a versatile, wearable, fully chip-integrated, nanofluidic interface incorporating field-effect transistor (FET) sensors. This will be the first commercially available chip of its kind in a wearable format for lab-on-skin applications.

Sweat has only recently emerged as a truly valuable means of assessing key health data in real time, without resorting to invasive methods. We do not believe blood testing could or should be replaced: instead, we believe that sweat testing can and will become a powerful complement to blood testing, to monitor our health discreetly and to alert us when key signals start to behave differently. Xsensio is targeting a number of health and wellness applications around its unique Lab-on-Skin™ sensing platform, ranging from skincare sensitivities and sports-related electrolyte imbalances to cystic fibrosis and hypertension.

FinFET sensing chip

Xsensio positions itself as a business-to-business (B2B) company, undertaking research and development and commercial partnerships with large consumer care and consumer electronics companies. Over time, Xsensio will expand its portfolio of B2B partners to include medical technology and drug companies, in order to offer companion diagnostics tools for therapeutics.

TETRACOM: invaluable support for tech transfer

The TETRACOM grant has allowed Xsensio to achieve two goals: first, Xsensio will soon be able to add a key component to its final Lab-on-Skin™ sensing platform solution to make that solution truly wearable. Second, Xsensio will be able to capture the intellectual property (IP) arising out of the on-chip integration of the nanofluidics together with the sensors.

> *"TETRACOM provides SMEs and start-ups with a clear framework for collaboration with university labs"*

Often, highly promising technologies do not progress beyond the invention stage, never becoming true innovations that disrupt the marketplace. The reasons are numerous, but the cost of making that transfer can be exorbitant, especially for small/medium enterprises and start-ups that cannot afford to allocate the time and resources necessary to work with a university lab. A programme such as TETRACOM is therefore excellent as it provides SMEs and start-ups with a clear framework for collaboration with university labs, one that provides milestones and deadlines and a lock-in of IP at the end of the collaboration.

> *"Xsensio is targeting health and wellness applications from skincare sensitivities and sports-related electrolyte imbalances to cystic fibrosis and hypertension"*

A single-board computer made in Europe

Cyber-physical systems meet supercomputing

In April, we saw another indicator of the booming popularity of the DIY electronics scene, when the Kickstarter campaign for the UDOO X86 board smashed its €100,000 target overnight. Thanks to the EU-funded AXIOM (www.axiom-project.eu) and Mont-Blanc (www.montblanc-project.eu) projects, a new and improved version of Barcelona Supercomputing Center's OmpSs parallel programming model can be run on a cluster of UDOO X86, allowing hobbyists and professionals to craft their own supercomputer. The AXIOM team aims to create a single-board computer – a complete computer comprising microprocessor(s), memory, input/output and other features on one circuit board – which is designed and manufactured in Europe.

HiPEAC caught up with UDOO co-founder Maurizio Caporali (MC) of the University of Siena and Xavier Martorell (XM) of Barcelona Supercomputing Center to find out more.

What's so special about this new board?

MC: UDOO X86 is a unique single-board computer: it's both the world's most powerful maker board and a fully fledged Arduino 101. As a computer, UDOO X86 is a quantum leap forward compared to regular single-board computers for makers, and its performance is comparable to most notebooks. It can drive up to three 4k screens – that is, screens with a horizontal resolution of around 4,000 pixels – simultaneously and runs Windows (including Windows 10), Android and Linux. It is 10 times more powerful than the Raspberry Pi 3. Despite this incredible power, its Intel Quad Core 14nm 64-bit processors consume as little as 5-6W in energy, depending on the UDOO X86 model.

UDOO X86 has the same pinout as an Arduino 101 and is 100% compatible with Arduino shields, sensors and actuators. It can even run the Arduino integrated development environment directly from the main Intel quad core processor. The Arduino 101-compatible microcontroller is based on Intel Curie, which integrates 32-bit Quark SE system-on-chip, six-axis motion sensors and Bluetooth low energy. Last but not least, UDOO X86 is open source and open hardware.

UDOO X86: Vital statistics

- Processor up to 2.56GHz
- Up to 8GB of RAM
- Drives up to three 4K monitors simultaneously
- Completely Arduino 101-integrated
- Runs any X86 Linux distribution, Windows and Android
- Multiple options for mass storage
- Ability to start up processor through on-board microcontroller

Why are do-it-yourself (DIY) electronics so popular? What are the benefits of making things open source?

MC: Hardware is becoming less expensive year by year, and people have started realizing that they can build their own stuff instead of buying it. Recently there's also been more focus on STEAM (science, technology, engineering, arts and mathematics) fields. What we are witnessing is not just a bunch of hobbyists; it's a new industrial revolution, embodied by makers.

But there's another reason behind the success of DIY electronics: it's a natural process, driven by human curiosity. So far people have programmed pixels on a screen. It's a natural progression for those same programmers – as well as new generations – to program the world around them.

This is only possible thanks to the open source community, which thrives around these technologies. Releasing a technology under an open source licence makes it truly accessible, and is the best way to get traction and attract developers. We are proud to be 100%

open source: the schematics, binaries, source code, 3D files and the bill of materials for our boards are all available on our website.

What are some of the challenges for future cyber-physical systems (CPS)?

XM: CPS platforms are still less powerful than those of laptops. In these environments, the biggest challenge is supporting the needs of applications relating to computation and communication performance and low power dissipation. As a result, further research needs to be done in both hardware and software.

The hardware infrastructure for CPS is organized differently to current clusters, or the cloud. In CPS, nodes can be as small as an Arduino shield providing a sensor, or as big as a fully fledged board with high computational capabilities. But they have one characteristic in common: they should be able to generate or process communications at high bandwidth. CPS is another source of so-called 'big data', and we will need to learn how to process all the information generated.

In terms of the software ecosystem, in recent years there has been huge growth in support for standard programming languages and environments on embedded boards. The Mont-Blanc and AXIOM projects are going in this direction. Nevertheless, support for small devices and accelerators still needs significant investment to make them usable for the community.

Do we have the capacity in Europe to design and manufacture a single-board computer? What innovations does this board offer?

MC: AXIOM benefits from the expertise of SECO (www.seco.com), a company with over 35 years' experience in embedded electronics, which sells and ships its products all over the world. The two market uses we are focusing on are home automation and smart video surveillance.

The UDOO X86 AXIOM cluster

Using field-programmable gate arrays (FPGAs) is not new in this field. The revolution offered by this board is how easy it is to interconnect and program.

What is the role of OmpSs in this project?

XM: The OmpSs programming model aims to provide a style of programming which can be used across different applications to exploit parallelism and acceleration. With the cluster version, we support the execution of applications across several nodes without needing to explicitly indicate the communications that should take place, just by providing data directionality hints to the compiler. OmpSs allows certain tasks which take a long time on the microprocessor to be accelerated using FPGA devices, like signal processing algorithms, encryption, or even some number-crunching applications.

How can we exploit supercomputing technologies for CPS?

XM: In the old days, programming for small devices had to be done in assembly language. This is no longer the case: most devices now carry a relatively powerful microprocessor which can be used to run a standard compiler toolchain in order to generate efficient code. Current cards are powerful enough to sustain the compilation of the system by themselves, although some cross-compilation tools will still be useful. In the near future, we'll see less need for cross compilation, and platforms will become ready to run code that was initially designed to be run on a supercomputer.

With the support of existing libraries from supercomputing, we'll be able to improve the performance of both computation and communication, while reducing power consumption. For example, if the compiler is able to generate single instruction, multiple data (SIMD) code for the target architecture, energy consumption will be reduced due to better exploitation of hardware resources; and, at the same time, execution time will decrease. Also, current platforms allow parallelism to be exploited in the same way as in supercomputers, so most of the software is perfectly portable. Communications is perhaps the most challenging area of CPS, because the environments are less coupled, but the needs regarding data bandwidth are the same.

"The revolution of this board is how easy it is to interconnect and program"

New European research on image processing and heterogeneity

Recently launched Horizon 2020 projects will provide a reference platform for image processing applications and help developers exploit heterogeneous hardware.

TULIPP: TOWARDS UBIQUITOUS LOW-POWER IMAGE PROCESSING PLATFORMS

Many industrial domains rely on vision-based applications which require compliance with strict performance and embedded requirements. The EU Horizon 2020 project TULIPP will develop a reference platform (covering the hardware, operating system and tool chain), complete with implementation rules and interfaces, to tackle power consumption issues while delivering high, efficient and guaranteed computing performance for image processing applications.

The aim is to counteract the constant changes in image processing boards due to the continuous evolution of components. The overall design cost of image processing devices will be drastically reduced by developing and integrating universal interfaces on the platform, meaning that any new generation of components may be integrated without significant additional design costs.

To demonstrate its effectiveness, an application-driven version of the reference platform will developed during the project, implementing:

- Reference hardware architecture: a scalable low-power board.
- Low-power operating system and image processing libraries.
- Energy-aware tool chain.

This will lead to three proof-of-concept demonstrators across different application domains:

- Medical imaging: Radiation dose reduction in surgical X-ray systems through real-time image processing.
- Automotive vision-based driver assistance systems: Reliable, affordable and low-latency image processing with hard real-time constraints.
- Unmanned Aerial Vehicles (UAVs): More intelligence through an energy-efficient and small on-board controller.

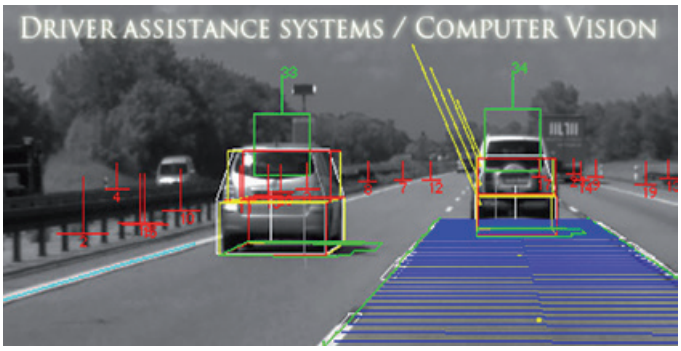

The TULIPP platform will be tested in three different domains

TULIPP will focus on design and implementation rules which will allow anyone to produce a compliant platform and benefit from the technological advances generated by the project. The interfaces will be defined so as to enable users to reproduce any subcomponent of the platform and plug it into existing platforms.

Using this reference platform will enable designers to develop an elementary board at reduced cost to meet typical embedded systems requirements relating to size, weight and power. Moreover, for less constrained systems whose performance requirements cannot be fulfilled by one board, the reference platform will also be scalable so that the resulting boards can be linked together for higher processing power. The TULIPP advisory board will supervise and monitor the progress of the guidelines

as well as testing the platform with its own applications and supporting the consortium in the standardization process.

NAME: TULIPP - Towards Ubiquitous Low-power Image Processing Platforms
START/END DATE: 01/02/2016 – 31/01/2019
KEY THEMES: embedded and low-power image processing, hardware platforms, multicore, real-time operating systems, heterogeneous architectures
COORDINATOR: Dr. Philippe Millet, Thales
PARTNERS: Germany: Ruhr-University Bochum, Fraunhofer IOSB; Norway: Norwegian University of Science and Technology; Sweden: Synective Labs; Belgium: Hipperos; UK: Sundance Multiprocessor Technology Ltd.; France: Efficient Innovation, Thales
BUDGET: €3.9M
WEBSITE: www.tulipp.eu
CONTACT: Contact@tulipp.eu
TULIPP is funded by the European Commission under the H2020 Framework Programme for Research and Innovation under grant agreement no 688403

SIMPLIFYING AND OPTIMIZING HETEROGENEITY WITH TANGO

Launched in January 2016 and coordinated by Atos, the TANGO project aims to help developers exploit the possibilities offered by customized heterogeneous hardware, simplifying the development process for applications in areas such as the internet of things, cyber-physical systems, wearables, big data and high-performance computing.

As the impact of heterogeneity on all computing tasks is rapidly increasing, innovative architectures, algorithms and specialized programming environments and tools are needed. TANGO focuses on designing more flexible software abstractions and improved system architectures to fully exploit the benefits of heterogeneous platforms, while addressing energy optimization at the same time.

The project will help control and abstract the underlying architectures, configurations and software systems while providing tools to optimize various aspects of software design and operations, including energy efficiency, performance, data movement and location, cost, time criticality, security and dependability on target architectures. Optimization tools include field-programmable gate arrays (FPGAs), application-specific instruction set processors (ASIPs), multiprocessor systems-on-chip, central and graphics processing unit chips and heterogeneous multiprocessor clusters.

Among the innovations to be delivered by TANGO are the following:

- Reference architecture and its implementation, to include the results of research in different optimization areas.
- Programming model with built-in support for various hardware architectures.
- New cross-layer programming approach for heterogeneous parallel hardware architectures, featuring (semi-) automatic code generation including software and hardware modelling.

In addition, TANGO will provide mechanisms to allow control of the heterogeneous parallel infrastructures. The most important results of the project will be released as open source. TANGO is also interested in joining forces with other research projects, initiatives and IT community organizations in a Research Alliance, to help nurture research collaboration, integration and the effective promotion of results.

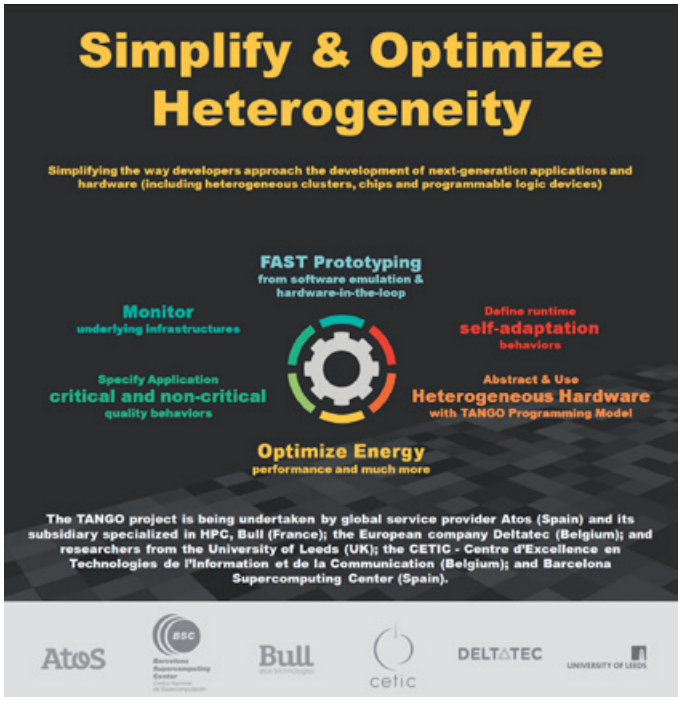

NAME: TANGO – Transparent heterogeneous hardware Architecture deployment for eNergy Gain In Operation
START/END DATE: 01/02/2016 – 31/01/2018
KEY THEMES: software engineering, operating systems, HPC, embedded computing, low-power computing, programming models, heterogeneous parallel architectures, security, privacy
COORDINATOR: Clara Pezuela, Atos
PARTNERS: Spain: Atos, Barcelona Supercomputing Center; France: Bull; Belgium: Deltatec, Centre d'Excellence en Technologies de l'Information et de la Communication; UK: University of Leeds
BUDGET: € 3.2 M
WEBSITE: www.tango-project.eu
TANGO is funded by the European Commission under the H2020 Framework Programme for Research and Innovation under grant agreement no 687584

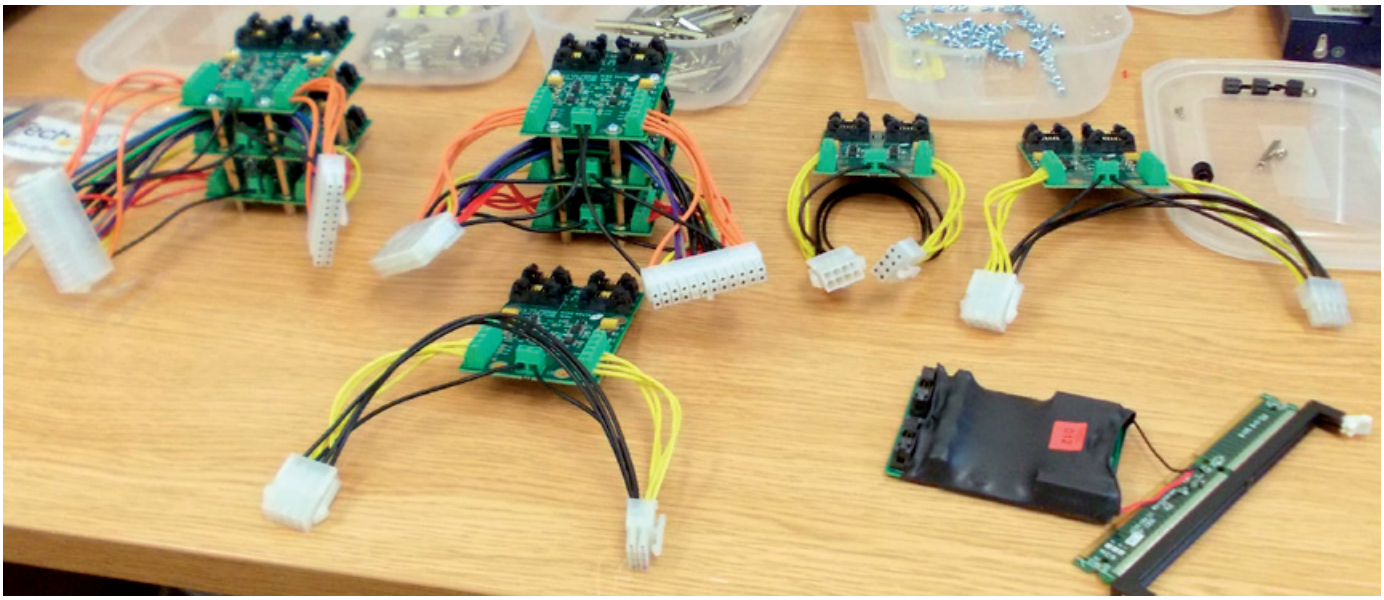

An Adept approach to power management

Mirren White, Edinburgh Parallel Computing Centre (EPCC), sets out how, through analysing energy consumption and power use in parallel software and hardware, Adept will help develop energy-efficient systems for the future.

Launched in September 2013, Adept is a three-year, European Commission-funded project working on bettering our understanding of energy and power use in parallel software and hardware, with the goal of using this knowledge to predict power use and energy efficiency in new systems. The project brings together experts from both the high-performance and embedded computing sectors, and utilizes their expertise to advance knowledge about the efficiency and power profiles of systems. This allows for the development of more economical, energy-efficient systems without the need for speculation.

Parallel computing is no longer limited to largescale HPC systems; in fact, parallel technologies are becoming critical to our everyday lives. Parallelism on every scale is on show throughout our societies, from the HPC machines in our labs to the smartphones in our pockets. Small and large businesses alike now need sensible, affordable parallel systems in order to remain competitive, and there is a vast array of different parallel commodity hardware now available. Investigating and increasing the efficiency of such devices is therefore no longer an abstract concern but a real and pressing need. Financial needs, environmental concerns, system requirements – all of these considerations, and more, will affect how systems are built in future.

This is where Adept comes in, by developing a range of tools to facilitate the efficient use of parallel systems.

The Adept Power Measurement System

One of the key outcomes of the project is our sophisticated Adept Power Measurement System (APMS). This fine-grained measurement infrastructure reads the current and voltage from the powerlines that feed the different components of a computer system, such as the. processor, memory or disk. The APMS is capable of measuring from multiple components with the very high resolution of one million samples per second.

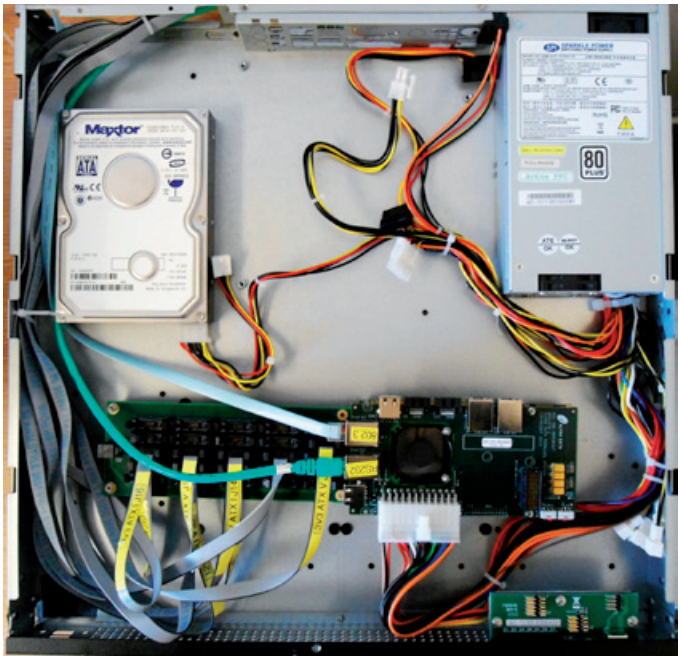

The APMS can provide detailed measurements from multiple components

Tracing earth-shaking performance and power

The APMS allows us to see in fine detail how software uses power on hardware over a given time. An example of this is our SEISMO case study. SEISMO is a code used to model seismic wave propagation through different materials; a typical application is in oil and gas exploration. Adept has simplified the original code, which works on a 3D domain, to compute over regular 2D grid and has created a number of different parallel implementations, such as MPI (message passing interface), OpenMP (open multi-processing), UPC (Unified Parallel C), CUDA and hardware description languages including VHDL, to support a range of different hardware architectures, such as central or graphics processing unit (CPU/GPU), or field-programmable gate array (FPGAs). We tested these new implementations on different systems to understand the impact on performance and power usage. The APMS gives us a very clear view of how the programme is using power throughout its run; an example of running SEISMO on an Intel Haswell i5 4670K desktop-class, quad-core processor, with a grid size of 5000x5000 elements over 10 iterations, is shown in the power trace below.

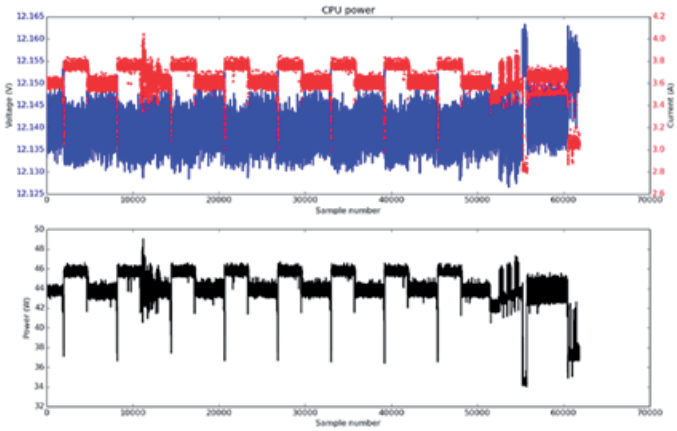

SEISMO power trace

The Adept Benchmark Suite

To complement the APMS, the Adept project has also developed a suite of benchmarks that can be used to test and evaluate existing systems. The benchmarks are designed to be used for system characterization and target specific operations and common computational patterns.

The suite consists of three different types of benchmarks:

- Micro benchmarks: small, single-purpose functions such as basic operations on scalar data types, branch and jump statements, function calls, input/output operations, inter-process communication, or memory access operations.
- Kernel benchmarks: computational patterns and kernels that largely consist of the operations from the micro benchmarks.
- Application benchmarks: small applications that consist of multiple computational kernels.

As part of our Benchmark Suite, we provide a wrapper for Intel's Running Average Power Limit (RAPL) system, which is an in-band method for reading power and energy counters on certain Intel CPUs.

Together, the Power Measurement System and Benchmark Suite form a powerful set of diagnostic tools to allow in-depth analysis of an application's power use in every aspect of a system.

The Adept Performance and Power Prediction Tool

Our work within the Adept project however is not limited to measuring power consumption. Another important outcome of the project is our Performance and Power Prediction Tool. Using detailed statistical modelling that examines a software binary, we are able to predict how well a CPU and memory hierarchy system will perform and how power efficient it will be, even if we do not have access to that system or even if that system does not yet exist. The Adept tool will be an asset to software developers and system designers, by freeing them from making poorly informed decisions about how to implement changes to their systems. The Adept tool allows for the design of smarter, cheaper, and more efficient systems, because a system's performance and power behaviours can be matched to a specific workload. Giving owners and developers the freedom and flexibility to know how their equipment will perform prior to porting their workloads means giving them the ability to make better choices about what they implement, how, and when.

In the final few months of the project, we will be focusing on improving the Adept tool wherever possible to makes its prediction increasingly accurate, and we will use our measurement infrastructure to conduct a wide range of experiments around power efficiency techniques in software development. But we will also focus significant effort on the exploitation of the project outcomes to ensure the lasting impact of our research. A lot of challenging work remains to be done, however we are certain that Adept will deliver on all its objectives and more.

Codeplay: the programming behind your car's eyes on the road

In the latest in our series on thriving small and medium enterprises (SMEs), Andrew Richards gives us an insight into how Codeplay responds to the challenges of delivering machine intelligence into embedded devices while keeping power use low.

COMPANY: Codeplay Software Limited
MAIN BUSINESS: Heterogeneous software platforms and tools, based on open standards
LOCATION: Edinburgh, UK and Toronto, Canada
WEBSITE: www.codeplay.com

Founded in Edinburgh by our chief technical officer, Dr Jens-Uwe Dolinsky, and myself, Codeplay initially developed optimizing compilers aimed at games programmers. Today, the company employs around 60 people and is internationally recognized for its expertise in advanced heterogeneous computing technologies, compilers and development tools.

Our growth today is driven by machine vision and artificial intelligence, which are now able to operate at reasonable performance levels for real-world applications. This is driving the revolution currently taking place in the automotive sector, with vision processing providing deep intelligence for advanced driver assistance systems (ADAS) and, ultimately, autonomous vehicles. To bring widespread embedded intelligence to market requires open standard software platforms that can both support the development of complex, highly parallel software, while also giving high performance on low-power parallel processors. As we move into safety-critical or regulated markets, we also need to ensure that the toolsets deliver the assurances required of software in those sectors.

Codeplay provides the backend development necessary for autonomous cars Photo: © Haiyin | Dreamstime.com

OpenCL (Open Computing Language), The Khronos Group's standard for parallel programming of heterogeneous systems. is now being adopted within the automotive sector and cloud data processing. OpenCL enables complex algorithms to target specialized accelerator processors with a standard programming language, while promoting reusability and cross-processor programming.

Codeplay provide leadership in Khronos Group standards to ensure that our software provides best-in-class performance with ease of integration. ComputeCpp enables easy integration of C++ applications into complex heterogeneous compute systems, while ComputeAorta is at the heart of our compute technology, implementing OpenCL, SPIR (Standard Portable Intermediate Representation), HSA (Heterogeneous System Architecture) and Vulkan.

Our day-to-day work involves implementing OpenCL compilers and runtimes on parallel processors, that is, the low-level drivers that manage hardware resources and the runtime application programming interfaces (APIs) that integrate it all together. To help developers, we also build debuggers for those architectures, and lead the definition and implementation of HSA (Heterogeneous System Architecture) standards for system-wide debugging and profiling.

We've also worked on the SYCL for OpenCL standard, enabling more complex, performance portable software (such as machine learning) to be accelerated by a wide range of parallel hardware. By promoting SYCL ideas to the International Organization for Standardization C++ body, we hope that more developers can take full advantage of massive parallelism with great power efficiency.

The company has won a host of awards, including the prestigious Investors in People standard, the Elektra Research and Design Award and the Scottish Enterprise Award. In 2013, Codeplay formed a research and development division, which focuses on high-level programming models and development tools for heterogeneous systems.

DO YOU WORK FOR AN INNOVATIVE TECHNOLOGY SME?
Contact communication@hipeac.net with your story.

Critical systems in avionics: sky-high

Photo credit: master films/P. Pigeyre

Multicore systems offer many advantages - such as energy and cost savings – but can they be designed to cope with the safety demands of the avionics sector? Sascha Uhrig, Dietmar Geiger, Nicolas Valot, Jean-Claude Laperche, Serge Barbagelata and Denis Chapuis of the Airbus Group set out the challenges ahead.

New processing opportunities provided by multicores enable novel aircraft functionalities which will not only further improve safety and availability but also reduce fuel consumption due to lower weight and smaller space requirements, hence increasing eco-efficiency.

Components involved in cockpit functions are, in general, more safety critical than those in the cabin
Photo credit: master films/A. Tchaikovski

Functions in the avionic domain are assigned a certain Function Design Assurance Level (FDAL) according to their criticality with respect to safety. This ranges from A to E, with FDAL-A being the most critical level, while FDAL-E indicates that there is no impact on safety. Specific Item Design Assurance Levels (IDALs) for hardware devices and software components are derived from the FDAL. To get a rough impression of IDAL assignment, most components involved in functions of the cockpit and flight systems are rated as FDAL-A, -B, and partly -C while cabin systems are mostly FDAL-C, -D, -E, although they can also be FDAL-B and -A.

An example of a function is braking on the ground, which is highly critical: if it fails, the aircraft cannot stop at the end of the runway, with catastrophic consequences. This function can be performed using three different techniques (sub-functions): wheel braking, reverse thrust, and air braking. The first step is to evaluate the criticality of a failure on the part of the individual braking systems (functions) and assign an FDAL.

The criticality of failures may depend on the flight phase and on the time in which a failure is detected (or not). For example, an undetected failure in the wheel braking system may have catastrophic consequences, whereas a known failure in the same system may have only major consequences, such as the crew choosing a longer runway and using reverse thrust and spoilers to decelerate.

Next, the systems are broken down into individual subsystems and, finally, into hardware and software components called items. These items are assigned individual Item Design Assurance Levels (IDALs) according to the FDAL and assignment regulations. By choosing specific architectures, i.e. control and monitoring

challenges for multicore development

paths, the IDAL of specific items can be reduced in some cases. If there is more than one item, independent items (those with no parts in common) need to be at least on the level below the FDAL. This means the items either need to be completely independent or one IDAL must be the same as the FDAL. Note that this simplification is just for demonstration; the actual regulations are much more complex.

To make things even more challenging, it should be noted that most of the above-mentioned items are also required for other functions; for example, air brakes are also used to reduce speed while the aircraft is in flight. Accordingly, the air brakes participate in the assignment processes of at least two functions. As a result, the assignment process of all IDALs within an aircraft forms a complex optimization problem, with safety and cost being target parameters.

Moving onto multicore processors, it's important to understand that the IDAL relates to the design of an item, i.e. how an item must be developed and structured such that the probability of malfunction or dysfunction is acceptably low. All IDALs, except IDAL-E, are critical and any fault should be detected and mitigated. Regardless of regulations, the target of all aircraft manufacturers is to provide reliable systems and to avoid any kind of incident.

Mapped to multicores, this means that any kinds of interference between items, in this case software components, need to be taken into account during assignment of the IDAL and/or they must be treated accordingly during execution to guarantee (virtual) independence. Otherwise the items' independence can no longer be guaranteed and the previously assigned IDALs are invalid.

Certification authorities monitor correct FDAL/IDAL assignment as well as the verification of correct implementation. In addition to respective standards, Certification Review Items (CRIs) define rules relating to certification of avionics systems, including multicore topics. Acceptable and unacceptable interferences between items on a single multicore are indicated here. Again, it should be noted that the aircraft manufacturer's interest in providing reliable products may be even higher than that of the certification authority.

Nevertheless, these products need to be as cost and energy efficient as possible and state-of-the-art technologies like multicore systems should also be used for mixed-critical systems. Since commercial-off-the-shelf (COTS) multicores are not designed with aeronautical standards in mind, they do not directly fit the requirements of the avionics sector and pose new challenges in avionics hardware and software design.

Photo credit: master films/A. Tchaikovski

An example of such challenges originates in the non-deterministic timing behaviour that can be observed in individual processor cores. Techniques to deal with, or, even better, circumvent this non-determinism need to be integrated into multicore avionics systems. Possible incarnations of such techniques could address item layout and mapping to cores, as well as suitable control mechanisms that guarantee the correct and timely behaviour of all items implemented by the multicore. Note that all software items need to fulfil the assumptions made during IDAL assignment, including their independence.

For safety reasons, the above described safety-critical items must feature the highest possible availability and correct functioning. In contrast, mission-critical items provide more flexibility. This is because a failure in a mission-critical item can lead to significant financial losses but will never result in human lives being put in danger. Mission-critical systems have many similarities with safety-critical systems, such as being based on the same design assurance principles, but they also have fundamental differences. For example, while a sensor in a safety system must never deliver incorrect data, i.e. it has to be shut-down and replaced by a redundant alternative sensor if a malfunction occurs, a sensor in a mission-critical system always needs to deliver best-effort data. The overall idea is to complete a mission in the best possible way.

Mapped to multicores, such mission-critical items can leverage the high multicore performance much better than safety-critical items, as long as all of them can fulfil their mission. Nevertheless, first single function IDAL-A multicore systems are already available; it will be just a matter of time until the first multicore is safely applied in systems with different IDALs and mixed safety and mission criticality.

Jan Henrik Weinstock and Rainer Leupers, Institute for Communication Technologies and Embedded Systems (ICE), RWTH Aachen University

Boosting the performance of virtual platforms using SystemC-Link

The use of virtual platforms (VPs) has become widespread among developers of embedded systems. Compared to real hardware platforms or field-programmable gate array (FPGA)-based prototypes, VPs offer optimal visibility into and control of the target hardware, enabling shorter architecture optimization iterations and more powerful embedded software debugging. Using VPs, engineers are well equipped to deal with the ever-increasing complexity of next-generation systems driven by Moore's Law.

In contrast to real hardware, however, the performance of VPs takes a hit when more processors are added into the design. The reason for this lies in the fact that the current standard for VP design – SystemC – only employs a single host processor for simulation. When moving towards VPs for multi- and many-core designs, this traditional, sequential simulation technique clearly becomes a performance bottleneck, threatening the use of VPs as a productivity tool.

VPs use simulation time based on the start of an event, as opposed to a real-time clock, which advances in a linear fashion. A typical SystemC-based VP sets a predefined quantum, that is, a portion of simulation time, for a specific number of simulation steps. If it has multiple processors, each one will have to execute its instructions before the next gets its turn. Once all the processors have executed their instructions, the current quantum is completed and the simulation time is advanced to the next quantum.

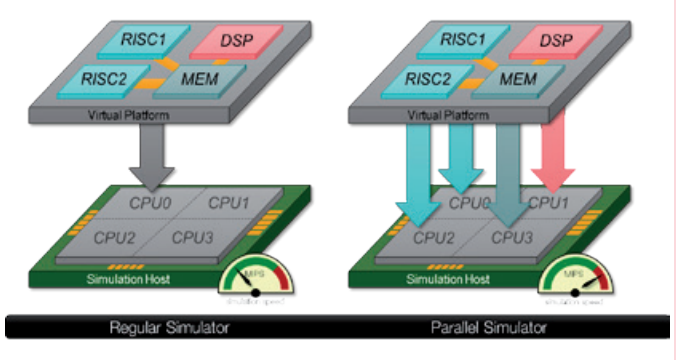

Conventional SystemC simulation environments only make use of a single host core even when simulating multiple guest processors. SystemC-Link allows simulation of different guest processors using all host processors

Conceptually, this leaves plenty of room for performance gains using parallel simulation, yet no approach has been adopted by the industry. One major reason for this is that many key players in the industry only provide models which are not safe for parallel execution. In order to overcome this situation, researchers at RWTH Aachen have developed SystemC-Link, a parallel SystemC simulation framework with the ability to emulate a sequential environment for models requiring thread safety, while utilizing parallel simulation to boost performance on multicore hosts.

In SystemC-Link, a virtual platform is comprised of segments. Each segment consists of one or more possibly interconnected simulation models and a copy of the SystemC simulation kernel. For models requiring augmented kernels or specific versions, the standard kernel can be transparently exchanged. Communication between models from different segments is carried out using connector blocks, which are placed along borders between segments. Segments form a virtual sequential environment for the models within, with connector blocks as the designated means to safely exchange data with the rest of the simulation.

A SystemC-Link simulation controller handles elaboration and simulation of those segments. Segments are simulated in parallel using an asynchronous scheme: local times of each segment are allowed to differ, as long as they remain within a sliding temporal window, defined by the slowest segment and the length of the quantum. Having a time window gives the simulation controller a higher degree of freedom when selecting segments for simulation, which in turn leads to more activity that can be handled in parallel and ultimately to simulation speedup.

The SystemC-Link approach has been tested with multiple VPs from different hardware architectures, including the distributed and tiled architecture EURETILE as well as an OpenRISC-based SMP MPSoC design. Both platforms feature a configurable number of processors for simulation (64-512 for EURETILE, 1-8 for OpenRISC) which naturally tailor well to parallel simulation approaches. Using an off-the-shelf desktop computer with four processor cores – a typical specification for a workplace PC – SystemC-Link achieves performance improvements on real-life benchmarks of between 3.2 and 3.7 times for both platforms compared to standard SystemC simulators. Consequently, these

parallel simulators increase productivity by providing early results during design space exploration and shortened debugging cycles.

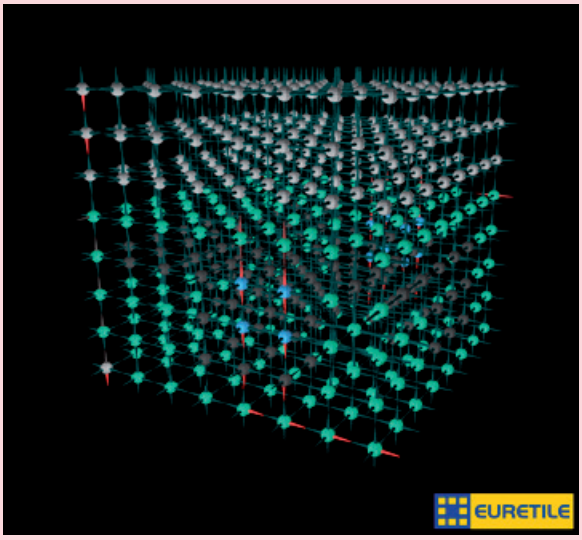

The EURETILE virtual platform using parallel SystemC technology. Each one of the 512 spheres represents a computational tile equipped with a RISC processor, memory and a network processor, which interconnects tiles in a toroidal 3D grid (red interconnects illustrate communication activity)

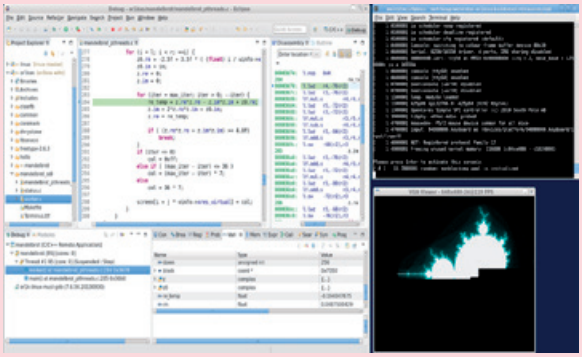

OpenRISC SMP virtual platform with complex CPU and peripheral models capable of running the Linux kernel. Debugging of the target application is improved by faster simulation speed thanks to parallel SystemC simulation technology

FOR MORE INFORMATION ON PARALLEL SYSTEMC SIMULATION
visit the following URL or scan the QR code below
https://www.ice.rwth-aachen.de/research/tools-projects/entry/detail/parallel-systemc-simulation/

Technology opinion: The case for open source hardware

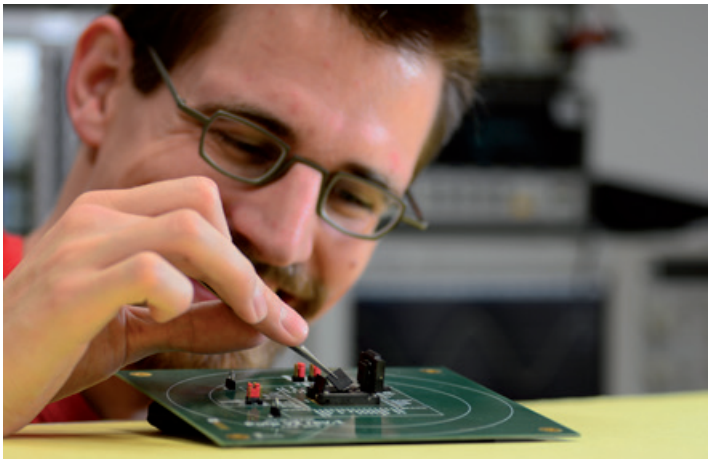

Frank K. Gürkaynak, ETH Zürich, and Luca Benini, University of Bologna

At ETH Zürich and the University of Bologna, we have been working on parallel, ultra-low power processing systems (PULP) using an open source approach. Our goal is to make all source code that we develop in this project freely available for others to use and build on. As a first step, we have released a silicon-proven single-core microcontroller called PULPino (http://pulp-platform.org), complete with support tools, that implements the similarly open RISC-V architecture (http://riscv.org) by Berkeley using a permissive Solderpad licence (http://solderpad.org/licenses) derived from the Apache v2.0 licence explicitly to support open source hardware.

So why did we decide to release our work using open source licences? For us, this was a strategic decision and we see many important benefits. The open nature of our project allows us great freedom in our collaborations with both industry and academia. There are no negotiations, limitations and delays involved when collaborating with an external partner which makes it much easier to start doing useful work. Our partners do not have to live with our design decisions; if there are things they would like to do differently, open source releases give them the freedom to do just that.

Conversely, in the long run we also hope to benefit from contributions from other groups that use our platform. Open hardware platforms are great for benchmarking as everyone will have access to exactly the same environment. Finally, usable open source hardware will lower the entry costs for integrated circuit design for some SMEs

and allow them to be more competitive, hopefully creating more business and employment opportunities.

For digital systems, a hardware description language (HDL) is usually provided, setting out the structure and behaviour of electronic circuits. It is this description – in our case, SystemVerilog – that we are currently releasing as open source. Most chip designers would agree that there is much more to a good usable design than its HDL description; as documentation and verification methods are just as important, we also include these in our releases. The other point that potential users look for is whether or not the design was actually manufactured and tested. We have to date taped out more than 15 chips as part of our PULP project (http://asic.ethz.ch/applications/pulp.html) using a number of different technologies, which was a very important factor for some of the groups that have decided to use our system.

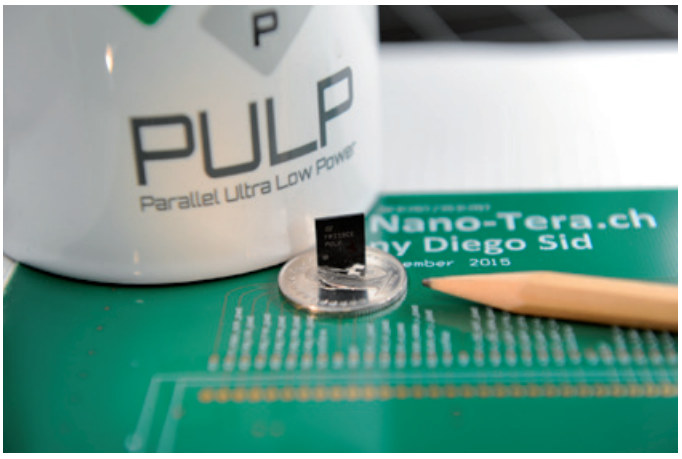

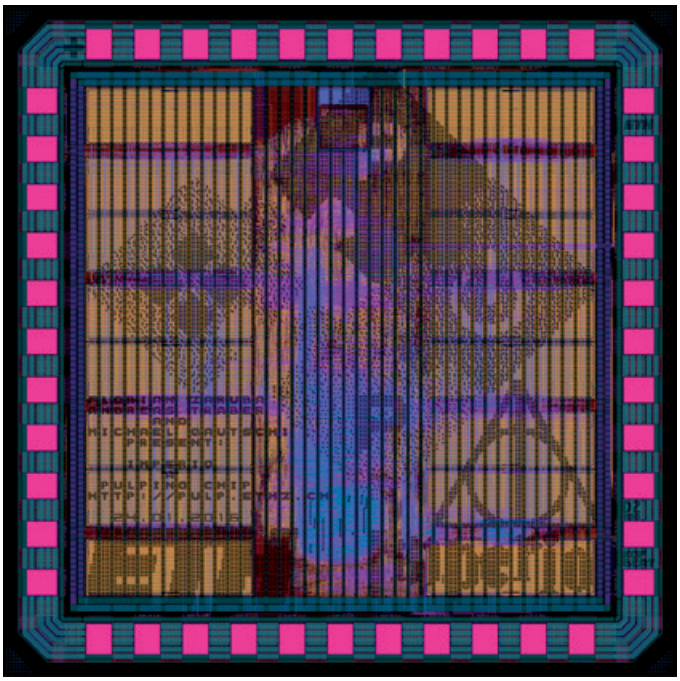

More than 15 PULP chips have been taped out (above). Below: PULPino chip layout

Of course, if someone were interested in making their own chip using PULP, they would have to work with design implementation tools that need gate libraries, timing and power models, for example, which are proprietary and not open-source. Interestingly, the situation is different when the same HDL description is used for field-programmable gate array (FPGA) designs. The HDL code will have to be mapped onto the resources of the specific FPGA by using steps similar to an application-specific integrated circuit (ASIC) flow. However, FPGA vendors have an interest in selling FPGAs, so do not tend to object to making the final configuration of the FPGA with our design freely available. After all, this configuration (typically called a bitfile) will only work if you purchase the specific FPGA from the vendor. In fact, we have released the bitfile of our PULPino system for the popular ZED board using the Xilinx ZYNQ series of FPGAs.

Open source software has been immensely successful and has helped to create enormous value for science and industry. In some ways, open hardware is still where open source software was about 20 years ago: not yet there, but seriously knocking on the door. We are by no means the first group to release hardware in the open source domain, but we are putting significant effort and resources into releasing high-quality, silicon-proven and competitive designs.

Second, we are releasing a complete package, including the support infrastructure (compilers, debuggers, verification scripts) which are essential to work with the hardware. Third, this is not a single-shot release: what is available today is sufficient to build your own single-core ultra-low power 32-bit microcontroller system-on-chip, while the next releases will open the way for multicore systems, with many advanced features supporting near-sensor processing for the internet of things.

Finally, we are releasing this under a permissive licence which allows closed-source derivatives and commercial exploitation. We believe that the open source hardware movement will play an important role in the future in both industrial and academic settings, and we are very excited to be at the forefront.

"Open hardware is where open source software was about 20 years ago"

Career talk: Mafijul Islam, Embedded Software Manager, Volvo AB

Describe your role briefly.
I currently have two roles in the area of embedded software and engineering methodologies in the Advanced Technology and Research division of Volvo AB. One is leading and facilitating a group of research and software engineers, and the other is preparing and driving forward the Technology Development (Research and Advanced Engineering) Portfolio for Volvo Group Trucks Technology.

In the first role, I lead a team of around 15 people who have, on average, over 12 years' experience as research and software engineers in automotive embedded systems. This role includes supporting team members not just in their day-to-day activities but also in moving forward with their career paths and in achieving the goals of the Volvo Group.

The second role involves preparing the technology development portfolio, taking short-, mid- and long-term challenges into account. This requires an understanding of the market and technology trends, as well as customer and product needs. It also entails plenty of discussion with relevant stakeholders within the company to identify their needs.

How did you come to be interested in this area? What did you study at university and what training did you undertake subsequently?
I studied at Chalmers University of Technology in Gothenburg, Sweden, obtaining an MSc in computer engineering with a focus on dependable computer systems before getting my computer engineering PhD, specializing in computer architecture. Chalmers pursues research activities in close collaboration with Swedish industries, including Volvo AB. As a result, I was able to find out about current research and development industries within the automotive industry during my studies, which inspired me to start my career as a research engineer in this sector.

While working as a research engineer, I took an active part in several areas of technology research in the context of embedded systems, participating in a number of publicly funded research projects comprising academic and industry partners from Sweden

and the wider European Union. I also worked on several internal advanced engineering projects with highly experienced colleagues. Participation and collaboration in such projects, in addition to in-house training and the working culture at Volvo, have been crucial in helping me acquire the skills needed for my current work.

What's a typical day like in your job? What aspects do you enjoy most?

My day-to-day work mostly involves meeting people, prioritizing action items, taking decisions (both administrative and technological), following up on activities, supporting group members and keeping my eyes open for emerging technologies. Tasks include speaking to a range of stakeholders to understand and prioritize their needs; preparing, leading and following up on technology development projects; and setting up goals for group members. A key feature of my work is communication, both within Volvo and with external stakeholders. I need to think not only about the physical working environment, but also emerging technologies, such as automation and autonomous vehicles.

Aspects which I find most enjoyable include meeting people, discussing different issues with different stakeholders, supporting people in their daily work and, above all, learning new technologies. I enjoy the variation and abstractness involved in my work.

How does working in industry differ from research?

Industry-based research is usually more applied than basic, more business driven and less publication driven. Also, working in industry is a combination and delicate trade off of short-, medium- and long-term goals and needs. Setting priorities is integral to our day-to-day work; however, priorities are not always static and can change drastically due to market and technology trends. In the case of the automotive industry, it's relatively new and very fast changing in relation to other industry areas, especially in terms of information and communication technologies, for example. As a result, we not only need to go deeper into certain areas but also to gain a broad overview.

> *"Industry-based research is usually more applied than basic, more business driven and less publication driven"*

In industry, as we produce products and services to meet the needs of end users, our tasks are necessarily customer driven. As a result, quality, safety and related issues are high on the agenda in the automotive industry. There is an urgent need to commercialize research activities within a given timeframe and budget to offer benefits for customers and to secure competitive advantage. Complexity in both technological and organizational aspects can potentially be unique with respect to the basic research. There is a strong trend towards teamwork, collaboration and an agile way of working; while we have flexibility and independence in our individual tasks, it may be to a lesser degree with respect to academia. On the other hand, we can afford to have varied assignments to broaden skills and competences.

What key skills and experience are essential for this job? What kind of personality do you think is best suited to it?

The two roles, general management and technology management, require a diverse set of skills. Essential skills include a good understanding of the technologies relating to automotive embedded systems; a good understanding of business needs, as well as market and technology trends; and strong networking and communication skills both within and outside the company. To lead people, it's very important to develop and maintain excellent working relationships with colleagues based on mutual respect and trust, in addition to being supportive and empathetic. In technology management, it's important to establish trade-offs taking multiple constraints into account and prioritize the right technology, at the right time, with the right people on board.

Mafijul works on emerging technologies at Volvo, such as autonomous vehicles. Photo © AB Volvo

WHERE WILL YOUR CAREER TAKE YOU NEXT?

Check out the numerous job opportunities on the HiPEAC jobs portal: www.hipeac.net/jobs

If you're passionate about your career and would like to share it with the HiPEAC community, we'd love to hear from you. Email communication@hipeac.net with your story.

HiPEAC internships are a great way to get a funded work experience placement at a great technology company – they can even lead to a job offer, as in this example. Interested in applying? Further information is available on the HiPEAC website: www.hipeac.net/mobility/internships

HiPEAC internships: your career starts here

FROM INTERN TO SENIOR RESEARCH ENGINEER AT SAMSUNG ELECTRONICS

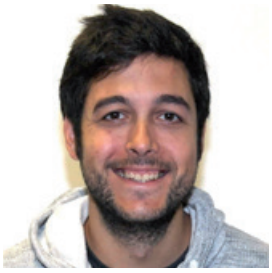

NAME: Albert Saà Garriga
RESEARCH CENTRE: Universitat Autònoma de Barcelona
HOST COMPANY: Samsung Electronics Research Institute UK
DATE OF INTERNSHIP: 31/08/2015-18/12/2015

The focus of my doctoral research was a set of tools exploiting parallel and heterogeneous computing to speed up and increase the energy efficiency of sequential source codes through source-to-source compilers, which translate code from one programming language to another. Thanks to HiPEAC, I was given the opportunity to undertake a four-month internship at the Samsung Research Institute in the UK. During this time, I worked on creating and optimizing a mobile phone augmented-reality application, based on a real-time 3D reconstruction system using monocular cameras.

Confidentiality issues mean that I can't go into further detail on the application at this stage, but what I can say is that this experience allowed me to bring together all of the knowledge I've acquired from my BSc in computer science, continuing with an MSc in computer vision and culminating in my PhD. It was much more than four months of interesting work; it opened up a new door in my professional career, as I was offered the post of senior research engineer working at Samsung. I would like to thank HiPEAC for providing me with this opportunity and would recommend it to all PhD students.

"The internship was much more than interesting work – it opened a new door in my career"

Prashant Sharma, Technology Manager at Samsung Electronics, commented: 'Albert's internship helped Samsung to start a new research project in augmented reality on mobile devices. His knowledge and experience were very useful for Samsung's research goals. After completing his doctorate, he re-joined Samsung as a full-time employee, where he continues his work in the same field.'

HiPEAC internships at a glance

- At the beginning of each year, HiPEAC member companies are invited to submit research projects for which they are looking for intern.
- Students are then invited to submit an application for an internship. They can apply to a maximum of two companies and cannot apply for more than two internships within the same company.
- The companies decide which applicants are eligible for the available internship positions. Final allocations of grants are decided by the HiPEAC Steering Committee.
- Internships for small/medium enterprises (SMEs) are fully funded, while larger companies get 50% funding for internships.
- Internships typically cover a three-month period.

www.hipeac.net/mobility/internships

Collaboration grants allow PhD students and junior post-doctoral researchers within the HiPEAC network to work jointly with a new research group to work on key challenges for computing systems. For further information about how to apply, visit www.hipeac.net/mobility/collaborations.

Creating the future through international exchange: HiPEAC collaboration grants

NAME: Alexandra Kourfali
RESEARCH CENTRE: Ghent University
HOST INSTITUTION: European Space Agency
DATE OF COLLABORATION:
01/09/2015 - 03/12/2015

LIFT-OFF FOR RELIABLE SPACE SYSTEMS

Electronic systems used for space applications, such as satellites for space exploration, raise unique challenges for engineers. In terms of computer hardware, static random-access memory (SRAM) field-programmable gate arrays (FPGAs) are an attractive solution, as they provide high performance while allowing the device to be (re)programmed remotely. However, higher levels of radiation in these environments can cause permanent and non-permanent faults on the FPGAs' configuration bits.

Reassuring fault tolerance in computing systems that contain FPGA devices is the most important problem for mission-critical space components: it is crucial that faults can be detected swiftly and that the hardware can be reconfigured on the spot. Several mitigation techniques have been used in order to sustain the functionality of the design, where the faults are detected and corrected.

During my first three-month collaboration with the European Space Agency (ESA), I conducted research for a method to test a design with fault injection. We designed a technique that will allow us to implement parameterized hardware systems in which the parameters define different circuit instances, where each instance represents a change of state that is caused in a microelectronics device due to ionization from protons, otherwise known as Single Event Effects (SEEs). We then focused on designing an integrated fault mitigation scheme that adds an extra layer of reliability to commercial off-the-shelf (COTS) FPGAs that will allow them to be used safely in future space missions. This technique can add fault mitigation infrastructure in the design, in order to provide extra protection by performing run-time recovery just for targeted areas of the FPGA.

I then undertook at three-month follow-up at ESA in which we worked on the mitigation scheme. First, I built a simplified fault mitigation method of Triple Modular Redundancy (TMR) for basic designs, which is currently the most robust technique that ensures fault tolerance of a system. Next, we applied a technique that allows forced scrubbing periodically in targeted areas in the FPGA's reconfiguration memory that are prone to failures. This is called microscrubbing and it is based on a technique called microreconfigurations (created at Ghent University). With microscrubbing, the design is periodically specialized only for the specific bits that are susceptible to a SEE. The microscrubbing system was presented during Space FPGAs Users Workshop in the Netherlands, while the mitigation scheme with TMR and microscrubbing were presented at the Military and Aerospace Programmable Logic Devices Workshop, USA.

As well as allowing me to put my research into practice for future real space applications, this partnership extended my network of contacts and opened up new opportunities for the future. Special thanks to David Merodio Codinachs and the entire Microelectronics Section at ESA/ESTEC.

The HiPEAC network numbers over 800 PhD students who defend, on average, more than two theses a week. Students affiliated to HiPEAC members can apply for internships and collaboration grants, as well as participating in HiPEAC networking events and attending the annual summer school.

Three-minute thesis

NAME: Ivan Vidović
RESEARCH CENTRE: Josip Juraj Strossmayer University of Osijek, Faculty of Electrical Engineering Osijek
ADVISERS: Dr Željko Hocenski, Dr Robert Cupec
THESIS: Image-based crop row detection using global optimization methods

FEATURED RESEARCH: CROP DETECTION FOR SMARTER, MORE EFFICIENT AGRICULTURE

You might think of agriculture as a low-tech sector, but it's becoming increasingly smart. Many processes in agriculture, such as planting, fertilization, plant protection and harvesting, require high precision yet are highly repetitive. These are exactly the sort of tasks which lend themselves to automation, making human labour less intensive and delivering greater productivity.

"Agriculture is becoming increasingly smart. To automate processes, accurate crop row detection is required"

Image-based crop row detection methods can help automate agricultural processes

In order to automate these processes, accurate crop row detection is required to allow automatic machine guidance. My research centres on developing image-based crop row detection methods. The main aim is to develop a method which is:

- highly insensitive to the presence of weeds and shadows
- capable of detecting rows of different crop types at different stages of growth
- capable of detecting straight and curved crop rows
- insensitive to the number and spacing of crop rows

My thesis proposes two methods to achieve this. The first, CRDI (crop row detection based on incremental method for line detection), is based on an incremental method to search for an approximate, globally optimal partition of a set of data points and on the DIRECT algorithm for global optimization. The second involves template matching followed by global energy minimization, which we call TMGEM, using a dynamic programming technique. To allow accurate crop row detection, the global energy function combines image evidence and prior knowledge about the geometric structure of crop rows. TMGEM is insensitive to the number and spacing of crop rows and is capable of detecting rows of different crop types at different stages of growth; it is also capable of detecting curved crop rows.

The thesis also proposes a new evaluation framework consisting of a crop row image database, including images of maize, celery, potato, onion, sunflower and soybean, a manual ground truth image creation approach and two crop row detection performance measures. This enables new crop row detection methods to be compared with existing ones efficiently and objectively. Using experiments, the two methods were evaluated by comparing them with comparable methods, using performance measures. The results in both cases were positive: CRDI outperformed the other methods considered, while TMGEM significantly outperformed other methods in straight crop row detection and proved its capability of detecting curved crop rows.

HOW IS YOUR RESEARCH GOING TO CHANGE COMPUTING SYSTEMS (AND THE WORLD) FOR THE BETTER?
Email communication@hipeac.net with your three-minute thesis pitch.

HiPEAC
COMPILATION ARCHITECTURE

Network of experts in high-performance and embedded architecture and compilation

With over 1,600 members across Europe, HiPEAC is **one of the largest networks** of its kind in the world.

Essential research

"Whenever we want to know where technology is going and what this will mean, HiPEAC is the source of information."

Sandro D'Elia, European Commission

Powering industry with technology

"HiPEAC allows us to keep up to date with current research into areas of business interest and monitor future trends."

Glenn Farrall, Infineon Technologies UK

Career opportunities, highly skilled candidates

"If you're looking for skilled PhD engineers in processor design, system architecture, compilers and tools, look in HiPEAC first, the best ones are there."

Christian Bertin, STMicroelectronics

Networking and mobility activities

"HiPEAC brings a lot of different people together, one of the first important steps when trying to take research out of the classroom and into practice."

Krisztián Flautner, ARM

Join the community

@hipeac

hipeac.net/linkedin

hipeac.net

This project has received funding from the European Union's Horizon2020 research and innovation programme under grant agreement no. 687698