HiPEAC info 70

OCTOBER 2023

Carmen G. Almudéver on scaling up quantum systems

Maximizing mobility: The computing challenges

Building a diverse computing systems community

contents

10

Carmen G. Almudéver on scaling up quantum

12

How EU research leads to two-wheeled innovation

14

Last-mile logistics get help from a HeRo

18

HPC meets CPS: How AMPERE turbocharges road and rail

23

Clue Technologies: High-performance computing for aerospace

32

How I became a computer scientist, with Panagiota Fatourou

Spanning the compute continuum from edge to cloud, HiPEAC (High Performance, Edge And Cloud computing) is a network of over 2,000 world-class computing systems researchers, industry representatives and students. First established in 2004, the project is now in its seventh edition. HiPEAC7 focuses on networking and roadmapping activities: bringing the computing community together in Europe, exchanging ideas, building thriving European value chains and exploring the long-term vision for computing systems.

hipeac.net @hipeac / @hipeacjobs

hipeac.net/linkedin hipeac.net/tv

Funded by the European Union

The HiPEAC project has received funding from the European Union's Horizon Europe research and innovation funding programme under grant agreement number 101069836. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

Cover image: faraktinov on Adobe Stock

Design: www.magelaan.be

Editor: Madeleine Gray

The theme of this magazine is mobility – a topic which, in 2023, evokes mixed feelings. On the one hand, mobility creates opportunities: it allows people to work, to study, to travel, and to ship products to customers worldwide. Safe and affordable mobility is one of the cornerstones of modern society. Historically, the mobility revolution started at the beginning of the 19th century with the invention of the steam locomotive, which made it safe and convenient to travel outside the region where one was living. Before the 19th century, long-distance travel was complex and dangerous, and it is no coincidence that the word travel derives from the French 'travail', meaning hard labour.

Today, we no longer associate travel with hard labour, but with holidays, tourism and fun. Smartphones allow us to stay connected with home while travelling, and card payments allow us to travel cash-free. The falling cost of transportation per kilometre and increased convenience has democratized travel, as well as facilitating international trade. Economically, the mobility industry (including transportation, tourism, and construction of infrastructure) adds considerably to countries' gross domestic product (GDP).

On the other hand, it has it dark side too: yearly, a staggering 1.3 million people are killed in mobility-related accidents, while up to 50 million are injured, the majority in low- and middle-income countries. It is the leading cause of death in children and young adults aged five to 29. In addition to this massive societal cost, there is also a huge environmental cost: globally, mobility causes 24% of energy-related greenhouse-gas emissions, almost equally distributed between passenger and cargo transport. On top of these emissions, there is also air pollution; as an example, mobility is responsible for about a third of all air pollution in the United States. This does not include the embodied emissions and pollution in vehicles and mobility-related infrastructure (roads, airports, oil production, …).

Given its crucial role in modern society and the economy, mobility is not going to disappear; on the contrary, it will undoubtedly grow. Hence, we will have to make it safer, and seriously reduce its ecological footprint. The HiPEAC community is very well positioned to make mobility safer and more eco friendly. I hope that this magazine will inspire you.

Take care

Koen De Bosschere, HiPEAC coordinator

HiPEAC brings together computing continuum partners to tackle research impact

In September, a trio of gatherings was held by the HiPEAC7 consortium to help further the project's objectives in three areas: stimulating innovation through the creation of value networks, establishing research priorities for Europe, and coordinating research and activities between major computing stakeholders in Europe.

The first meeting initiated discussion on building value networks. This particular meeting focused on optimizing the resilience of supply chains with cloud, edge, and the internet of things (IoT), including a view to bringing parts of the value chain back into Europe. Contributions from expert participants included the need to actively manage ecosystem engagement and to incorporate users as co-creators. They also discussed new imperatives such as the need for circularity, as well as how industry use cases dictate the essential components of computing value chains. For those who would like to participate, a follow-up meeting will take place next year.

The second meeting focused on aligning the strategic research agendas and roadmaps of major partnerships in Europe. Participants included Max Lemke and Jan Komarek of the European Commission, as well as representatives from ETP4HPC, Destination Earth, AIOTI, ADRA, BDVA, SNS, NESSI, ECSO, Chips JU, and FIWARE.

Finally, the third meeting aimed to help align activities between major initiatives influencing directions in the computing sphere, with AENEAS, EPoSS, Inside IA, ETP4HPC, Destination Earth, EUCloudEdgeIoT, Eclipse Foundation and 5G involved as well as HiPEAC. It was agreed to create a joint platform to facilitate this work – watch this space for more information.

Arm goes public, Intel spins off FPGA business

The chip-design company Arm has launched an initial public offering (IPO), marking its entry onto the stock market. The company started trading in New York on 14 September with 95.5 million shares and priced its IPO at $51 a share.

Based in the UK, Arm, which has been involved in HiPEAC since the network's founding in 2004, was acquired by the Japanese company SoftBank in 2016, which retains control of around 90% of the company's shares. According to the Arm website, the company's technology is used in more than 250 billion chips, spanning a spectrum of devices from sensors to smartphones to supercomputers.

Meanwhile the chip giant Intel announced in October that it would spin off its Programmable Solutions Group (PSG). The standalone company will begin operations in January 2024, and the chief executive will be Sandra Rivera, who currently leads Intel's Data Center and AI Group (DCAI). Although Intel will initially provide support, the company confirmed that it intends to hold an IPO for PSG in the next two to three years. Intel first broke into the field-programmable gate array (FPGA) market with its acquisition of programmable-logic company Altera in 2015.

Semiconductor industry players join forces to invest in RISC-V company

In August, semiconductor industry players Robert Bosch GmbH, Infineon Technologies AG, Nordic Semiconductor, NXP® Semiconductors, and Qualcomm Technologies announced that they would jointly invest in a company aimed at accelerating the commercialization of future products based on the open-source RISC-V architecture. The idea is to provide a one-stop shop to enable compatible RISC-V based products, provide reference architectures, and help establish solutions that are widely used in the industry. While the initial application focus is to be automotive, the joint statement by the companies said that there would be an eventual expansion to include mobile and internet of things (IoT).

Delving into computing systems from HPC to the edge at ACACES 2023

'I arrived with a poster and a suitcase, but I leave with knowledge, motivation, and cherished connections' – ACACES 2023 attendeee Cristina Olmedilla

In July, HiPEAC's annual summer school – Advanced Computer Architecture and Compilation for High-performance Embedded Systems – returned to Fiuggi for another highly successful edition. In a testament to the summer school's enduring popularity, applications for the nineteenth edition far exceeded the number of places available. Over 200 students eventually joined the event in Fiuggi, Italy, where they were treated to exceptionally high-calibre courses.

Topics included RISC-V (Filippo Mantovani), in-memory computing (Reetuparna Das), neuromorphic computing (Yulia Sandamirskaya), quantum computing (Koen Bertels), hardware security (Francesco Regazzoni) and sustainability (Carole-Jean Wu). As usual, they spanned the full compute continuum, from high-performance computing (J. Nelson Amaral gave a dedicated HPC course) to cloud (Babak Falsafi) to edge devices on the internet of things (Michele Magno). Meanwhile, this year's entrepreneurship course by Devin Voorsanger was a rousing, step-by-step guide to crossing the valley of death and successfully launching a technology business.

It is perhaps a sign of how ingrained machine learning and artificial intelligence (AI) are becoming in the field that two courses – Michael O'Boyle on automating compilers and Daniel Jiménez on machine learning for microarchitectural prediction – focused on using AI as a tool, rather than the usual courses on how to optimize machine learning itself.

The courses were complemented by a technical keynote talk by Axel Nackaerts (imec) titled 'The next branch', which explored clean computer design from a whole-system perspective. The Monday night entrepreneurial keynote offered lessons from the experience of Codeplay founder Andrew Richards.

Once again, dedicated sessions offered students a space in which to grow professionally. The much-loved poster session was as busy as ever, with senior attendees providing feedback on student's work. As for this year's careers evening, a personal, heartfelt talk by Blanca Sabater explored issues of diversity, focusing on neurodiversity and gender. Blanca was later joined by Reetuparna Das, Yulia Sandamirskaya, Axel Nackaerts and Filippo Mantovani for a careers panel sharing inspirational experiences. See p.34 for our interview with Blanca.

A number of attendees commented that they left this year's ACACES reinvigorated, with fresh motivation to continue their computing systems research. At HiPEAC, the staff have already started work on next year's edition – watch this space for further information.

Videos from ACACES 2023, including keynote talks and the first day of lectures, are available on the HiPEAC YouTube channel, HiPEAC TV: bit.ly/ACACES23_videos

TwinSpace opens up possibilities for energy and resource efficiency

Rolf Meyer and Mladen Berekovic (both University of Lübeck), Alexander Lehnert and Marc Reichenbach (both University of Rostock)

Traditionally, the optimization of non-functional properties of software, such as execution time, memory requirements and, in particular, energy, has been ignored because the speed of hardware doubles every two years without having to optimize these aspects. The unoptimized software leads to chips consuming more energy due to the necessity to exploit the higher computing capacity permanently.

With a view to making digitalization more sustainable, the TwinSpace project has taken up this challenge and intends to serve as a kind of hinge between economic and ecological benefits in the field of digital technologies. TwinSpace is a joint research project between universities (University of Rostock, University of Augsburg, University of Lübeck), small / medium enterprises or 'SMEs' (emmtrix Technologies GmbH, Tensor Embedded GmbH, AbsInt Angewandte Informatik GmbH) and large enterprises (CARIAD SE, e:fs TechHub GmbH). The project, an investment of nearly €7 million into energy-efficient mobility solutions, is funded by the German Federal Ministry of Economics and Climate Protection.

This groundbreaking initiative aims at optimizing (safety-related) software for embedded systems with respect to resource requirements (e.g. central processing unit (CPUs), graphics processing units (GPUs), and runtime), by mapping the application specification onto a twin space (black-box shells). The abstraction of components and thus decoupling of hardware and software enables mutually aware co-optimization. For this purpose, automated, decreasingly abstract reverse engineering of compiled code is applied, combining code analysis, mechanistic modelling and machine-learning approaches. The code of the application is then simulated on different emulated hardware platforms and the resource requirements are measured. Once the appropriate hardware platform is selected, the code is recompiled in accordance.

By abstracting from functional software and introducing software twins, it is possible to analyse and map software components onto real or virtual hardware at an early stage. Accordingly, the optimization goal of energy savings can be incorporated into the development process right from the start. The great thing about this is that reverse engineering does not create any cybersecurity dangers, as only non-functional aspects are transferred.

Hence TwinSpace, the green-coding development tool with optimized resource management, revolutionizes software development in the (safety-relevant) mass embedded systems market. Here, both economic aspects, e.g. through more efficient development processes, and ecological aspects, e.g. through increased energy efficiency or CO_2 savings through smaller chip sizes, are optimized.

Apart from the reduced need for computing capacity and the reduction of the energy consumption of systems, TwinSpace has a further impact on sustainable aspects by helping to minimize resource and energy consumption. One aspect, for example, is the reduction in the size of chip areas, while achieving the same functional results, saving further energy and rare minerals. A chip requires about 3.75 kg of CO_2 to manufacture. TwinSpace plans to reduce the required chip area by around 10%, which equates to a 375g saving of CO_2. Chips are produced in the billions or even trillions annually. With one billion chips, this means a saving of 375,000 tons of CO_2 just during the production of the chips. In addition, the hardware can be designed to be more efficient, more energy saving, and smaller, which directly leads to operational cost savings.

Putting this into the context of the United Nations' Sustainable Development Goals (SDGs), defined in the 2030 Agenda, TwinSpace positively impacts economic growth (SDG 8), innovation (SDG 9), sustainable production / resource scarcity (SDG 12) and climate-change mitigation (SDG 13), mainly through more efficient embedded systems in mass markets.

This project is intended to become a beacon, first radiating into the automotive industry and then into other industries, thus developing spill-over effects. The synergy between digital technology and sustainability that this project seeks to achieve will inspire further innovations, making TwinSpace a game-changer in the world of technology and environmental responsibility.

Keep up to date on project news and connect with the researchers involved by following the LinkedIn page:
linkedin.com/company/twin-space

MachineWare announces new ARM processor simulation and SystemC profiling products, adds Windows support

Rainer Leupers, RWTH Aachen

With the launch of the SIM-A simulator and inSCight profiler, MachineWare GmbH, headquartered in Aachen, is rapidly expanding its virtual prototyping products portfolio towards simulation of new processors and more efficient SystemC modelling.

SIM-A, a novel instruction-set simulator for ARM A- and M-class architectures, complements the company's lauded SIM-V product for high-speed RISC-V simulation. Like SIM-V, SIM-A has been derived from MachineWare´s proprietary Fast Translator Library (FTL), a framework for quickly building ultra-fast functional processor simulators.

"SIM-A is a game-changer for software development, helping our customers seamlessly bring up their complex software stacks, ranging from embedded Zephyr RTOS to Linux, at lightning speed,' says Lukas Jünger, MachineWare managing director and co-founder. 'With interactive debugging capabilities and the same ability to execute identical software as real hardware, SIM-A empowers you to conquer the most challenging software development tasks with ease.'

Another new software tool product, inSCight, addresses the need for efficient SystemC models for virtual prototypes of complex systems-on-chip. InSCight is a cutting-edge profiler designed for SystemC-based virtual platforms, offering a solution for high-speed simulation. With InSCight, users can swiftly pinpoint and resolve performance bottlenecks, unlocking substantial performance enhancements. The tool is indispensable in the context of virtual platforms, which are often composed of diverse models from various vendors and different abstraction levels, where inSCight efficiently identifies slow-simulating models.

'InSCight was designed to tackle real-world challenges, its primary mission being to empower our customers to maximize the potential of their virtual platforms,' adds Lukas. 'The benefit is optimal simulation speed for software development, regression testing and interactive debugging.'

In parallel to these developments, MachineWare has ported its software products to a wider range of host platforms. As a result, in addition to Linux, the entire tool suite is now available for MS Windows platforms. This includes existing products such as SIM-V, the Virtual Components Modeling Library VCML, as well as QBox ('QEMU in a SystemC box').

MachineWare will showcase the new products for the first time at the Design and Verification Conference & Exhibition Europe (DVCon Europe), a leading technical event on systems, software, design, verification, validation and integration, taking place on 14-15 November 2023 in Munich. A highlight of this event will be the presentation of the Renesas R-Car Virtual Platform, an example of how virtual electronic control units (ECUs) provide an indispensable tool for advanced driver assistance (ADAS) software development.

FURTHER INFORMATION:
MachineWare website machineware.de

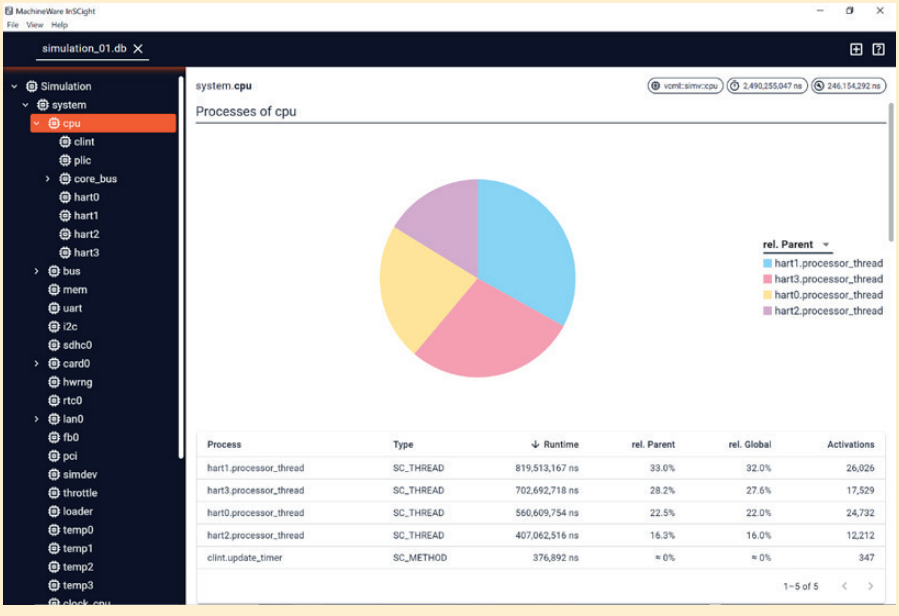

"SIM-A empowers you to conquer the most challenging software development tasks with ease"

Italian research centre propels innovation in HPC, big data and quantum

Marco Aldinucci, Università di Torino

Inaugurated in September 2022 with an investment of €320 million from the Italian state (thanks to European Union resiliency funds), the novel ICSC National Research Centre for High Performance Computing, Big Data and Quantum Computing carries out research and development for innovation in high-performance computing, simulations, and big-data analytics. The centre is organized into ten scientific areas, called 'spokes'. The first spoke is called 'FutureHPC' and revolves around next-generation hardware and software for HPC and the compute continuum (cloud, edge, etc.).

FutureHPC also unfolded in the inception of two national living labs, one in Bologna focused on 'Hardware and Systems' (HWS), led by Luca Benini, and one in Torino on 'Software and Integration' led by Marco Aldinucci. These labs aim at becoming lighthouses to serve as a cross-fertilization lab between academia and industry in the area of novel hardware, system software and tools for the compute continuum. They are also designed to cater for visiting researchers, providing visitors with an attractive co-working space with Luca Benini's and Marco Aldinucci's research groups and the laboratory equipment available for experimentation, starting from HWS RISC-V Monte Cimone cluster and SWI HPC4AI resources that include over 10,000 central processing units (CPUs) and over 120 graphics processing units (GPUs) that are fully dedicated to system software research.

FURTHER INFORMATION:

Spoke 1: FutureHPC and Big Data
supercomputing-icsc.it/spoke-1-future-hpc-big-data

HPC4AI
hpc4ai.unito.it

FutureHPC living lab on Software & Integration in Torino
hpc4ai.unito.it/livinglab-swi-icsc

Dates for your diary

ISC High Performance
12-16 May 2024, Hamburg, Germany
Workshop deadlines: 13 October and 7 December 2023
isc-hpc.com

IEEE Nordic Circuits and Systems Conference (NorCAS)
31 October – 1 November 2023, Aalborg, Denmark
vents.tuni.fi/norcas2023/

Chips Joint Undertaking launch event
30 November – 1 December 2023, Brussels, Belgium
bit.ly/ChipsJU_launch_event_news

HiPEAC 2024: High Performance, Edge And Cloud computing
17-19 January 2024, Munich, Germany
hipeac.net/conference
Sponsorship opportunities available
sponsorship@hipeac.net.

AEiC 2024: 28th Ada-Europe International Conference on Reliable Software Technologies
11-13 June 2024, Barcelona, Spain
Journal-track submission deadline: 15 January 2024
ada-europe.org/conference2024

HEART 2024: 14th International Symposium on Highly-Efficient Accelerators and Reconfigurable Technologies
19-21 June 2024, Porto, Portugal
Submission deadline: 18 March 2024
fe.up.pt/heart2024

Superposition, entanglement and the famous cat that refuses to die: the mysterious, winding world of quantum computing is both familiar and inscrutable. QUADRATURE, a new project funded by the highly competitive European Innovation Council Pathfinder programme, aims to bring scalable quantum computers one step closer to reality. We asked coordinator Carmen G. Almudéver (Universitat Politècnica de València) to guide us through this fascinating topic.

'The full potential of quantum is unknown, but its impact will be huge'

It feels like people have been talking about quantum computing for a while. When do you think it will become mainstream?

It's very difficult to say. Right now, quantum computing is at a comparable stage of development to classical computing 50 years ago, so we are still in the infancy of this technology. What I can say is that quantum computing has made remarkable progress over the last few years, from purely theoretical (quantum physics) to experimental processors running real algorithms. This progress has been steady and there have been breakthroughs every year, from demonstrating quantum advantage to scaling up systems to medium-sized intermediate scale, of a few hundred qubits. As an example, the IBM quantum roadmap foresees the release of a processor of more than 1,000 qubits by the end of 2023.

Most of these processors are available for cloud computing and they can run algorithms, although they are still very noisy and error prone. A higher number of qubits would allow us to correct some of these problems.

The main challenges in delivering workable quantum computers are scalability and fault tolerance. To be able to solve real-world problems, we need larger computers with thousands or even millions of qubits. To do that, we need to integrate more qubits on chips, but we also need to work on the different layers of the quantum computing stack.

The QUADRATURE consortium

It should be remembered that, in any event, quantum computers will not replace classical computers. Instead, they will extend the capabilities of classical computers, acting as accelerators for specific kinds of problem, in a similar way to artificial intelligence (AI) accelerators today.

What kind of problem could they solve, for example?

The potential of this technology is unknown: just as with AI, we don't know exactly how it will change our lives. That said, quantum computing is particularly suitable for optimization problems and for simulations, so we could see applications in materials science (such as creating stronger polymers), drug discovery, and logistics, for example.

Remind us what makes quantum computing so potent compared to classical computing…

So you know that classical computing works by using transistors to create binary values, bits, which are in a state of either 0 or 1, like a light switch that is on or off. The extraordinary thing about qubits is that they can be both states at the same time, like a spinning coin, a phenomenon known as superposition, which massively expands the computational capability. Qubits are implemented in different ways by different companies, with superconducting qubits and trapped ions being two of the most advanced in terms of qubit counts.

What about quantum communications?

Classical information transmission involves sending bits from part of the system to another using a classical channel. In contrast, quantum networks use entanglement to transfer information and need quantum communication links.

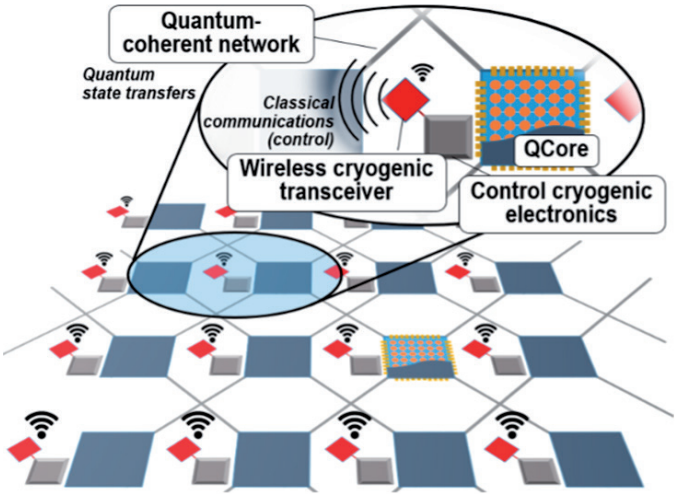

So how will QUADRATURE take us one step closer to quantum computing systems?

QUADRATURE addresses the scalability issue through an approach based on modular, distributed quantum computing technologies. Architectures composed of multiple cores (Qcores) will allow us to scale up quantum computing, just like the multicore concept in classical computing.

These Qcores will be connected via quantum-coherent links and via a cryogenic wireless interconnect, combining quantum and classical communication. This will allow us to link several quantum processors and so scale up the computer without having to integrate lots of qubits on a chip, which results in increased interference. The wireless networks, which draw upon concepts developed in the EU-funded WINC and WiPLASH projects, allow us to avoid overcrowding the architecture with wires.

Ah yes, we talked to Sergi Abadal about the wireless network-on-chip concept in* HiPEACinfo *67. This cryogenic business sounds pretty chilly , though…

Yes, the network has to work at temperatures of around 1-4K, as some qubit implementations have been demonstrated to operate at those temperatures. In QUADRATURE, we're aiming to scale up quantum computing systems by not only closing the vertical thermal gap, in which control electronic are placed near to quantum processors working at cryogenic electronics, but also the horizontal gap by using a modular multi-core architecture.

What are the main challenges you foresee in creating the QUADRATURE scalable, multi-chip architecture?

Interference is a major challenge. And, of course, this is an experimental project but also with an exploratory part performed by simulations. We're exploring quantum computing architectures from a full-stack perspective, from the technologies to the applications. As part of the project, we will build transceivers and quantum-coherent links, while simulation will be used to analyse the potential of this architecture and its ability to scale up. Simulation will also allow to see where we can optimize different aspects and to carry out gap analyses: if we do X, will we be able to get Y improvement for this application, for example.

Another issue is that qubits decohere over time, meaning that their quantum state is limited. So the timing in a quantum system is crucial.

It sounds like something out of science fiction: amazing, but hard to imagine in reality. Why should public administrations bet on high-risk technologies like quantum?

Of course, it's not certain that we will make it in quantum. But if we do, it will be a game changer: it has huge potential to solve problems that other technologies can't. This potential isn't just theoretical – it could have a massive impact on society and industry. For example, it could be used to better understand and combat climate change, or to further develop personalized medicine. Paired with other technologies, such as AI, it could become even more powerful. There's still a lot of work to do to get there, but we think the potential outweighs the risk.

FURTHER INFORMATION:
quadrature-project.eu

QUADRATURE has received funding from the European Union's Horizon Pathfinder programme, under grant agreement no. 101099697.

What do motorbikes have to do with the edge-to-cloud continuum, cybersecurity and homegrown European processors? HiPEAC caught up with Christos Kotselidis (PIERER Innovation) to learn how mutually beneficial exchanges of industry and academic knowledge contribute to powerful innovation.

Innovation powertrain

How research keeps PIERER Mobility in pole position

Headquartered in Austria, PIERER Innovation is a subsidiary of PIERER Mobility Group, the largest European manufacturer of powered two-wheelers (PTWs) – that is, motorcycles (KTM, Husqvarna, Gas Gas) and e-bikes (Husqvarna, Gas Gas, Felt) – and the fourth largest manufacturer of PTWs globally. The advanced research team, of which Christos forms a part, is tasked with investigating the long-term opportunities technology can offer mobility, with remarkable freedom to explore topics – from programming models, to security, to up-and-coming architectures like RISC-V, explains Christos. This knowledge is then transferred to the product groups for further development. 'The aim is to drive performance, quality, safety, and ultimately create a better user experience,' Christos notes.

One example of enhanced user experience is automation. Although self-driving features may not be the first thing that come to mind when you think of motorbikes, different riders have different expectations, as Christos explains. 'For track bikes, high performance and the driving experience are key,' he says. 'But after many hours on the motorway as part of a grand tour of Europe, for example, some assistance – such as self balancing or cruise control – may be very welcome. This has led to a trend to adapt advanced driving assistance systems (ADAS) to motorcycles, to help make riding more relaxing in specific situations.'

However, implementing ADAS on motorcycles raises major challenges. 'Space and weight are the main factors that limit the addition of computing capabilities. A modern car's capabilities are actually at the level of a small supercomputer – think of all the sensors collecting data, plus the ability to process this data in real time and make decisions. This is simply not possible on a motorcycle: every centimetre of space and every extra gram used has implications for the whole bike,' says Christos.

Another major challenge is the fact that a two-wheeled vehicle is constantly exposed to the weather elements, which means that the hardware has to be completely robust, Christos adds. 'Normally this comes with tight enclosures, which makes the computing system less capable of exploiting active cooling, meaning that you have to operate at lower frequencies to rely as much as possible on passive cooling. So you have to be frugal. In a car, you could overprovision the hardware and use it later to improve your use cases. In a motorcycle, it's the opposite: you have to use what you can fit on the bike, and then you have to optimize the code really hard in order to satisfy your use cases.'

Drawing on EU-funded research for industrial innovation

ELEGANT SecOPERA AERO

The research undertaken by the PIERER Innovation advanced research team centres around embedded hardware and software development, and this work is expanded by their participation in a number of projects funded by the European Union (EU), which help steer their vision. 'By taking part in research projects, we are able to acquire knowledge which could help us develop customized solutions (hardware and software) for our specific use cases,' says Christos. While striving for knowledge in itself is a

worthy goal, especially in the academic world, industry research tends to be constrained by the necessity for the results to have a swift impact on the company's bottom line. However, thanks to PIERER Innovation's forward-thinking, long-term vision, Christos says, the advanced research team is able to investigate topics which may not have an immediate effect but with the potential for major impact in the future.

The team is currently involved in three EU projects covering different aspects of the compute continuum: ELEGANT, SecOPERA and AERO. Funded under Horizon 2020 and now in its final year, ELEGANT is developing unified programming models for edge-to-cloud analytics, explains Christos.

'In the past, we had to develop code for our motorcycles and provide two implementations: one optimized for edge use cases, on a fleet of bikes for example, and the other a more relaxed, power-hungry version for the cloud. This means we have two versions of the same algorithm to process the same data. In ELEGANT, we're developing a new programming model where the user writes their software once, and then the system creates executables that can run either on the edge or on the cloud.' This reduces the amount of code to be written and therefore reduces costs, he notes. It also allows the team to dynamically control where the code will run and where the data is streamed, which is particularly useful for vehicles operating in different countries with different data-streaming tariffs.

With regulations such as the Radio Equipment Directive and EU Cyber Resilience Act, cybersecurity is another issue which has become a top priority for automotive manufacturers, explains Christos. 'The transition to secure operations affects every part of the pipeline, from specifications to manufacturing, sales, post-sales and even the retirement of the motorcycle.' This transition is the subject of SecOPERA, to which PIERER Innovation contributes the automotive use case. 'We're developing tools and methodologies (specifically for open-source software and hardware) along with verification processes to ensure cyber-security compliance throughout the whole product line. Since our parent company is an original equipment manufacturer (OEM), the aim is to demonstrate how the tools developed in this project can help us reach our cybersecurity goals,' he says.

Another major issue for automotive manufacturers, which was underlined by the COVID-19 pandemic, was that of supply-chain disruptions. For some years, the EU has been building up its domestic chip-building capabilities, most clearly embodied in the European Processor Initiative) (EPI), the EU-funded project to bring a low-power processor to market.

As part of the EPI and its affiliated projects, an automotive architecture is being developed, which caught the attention of the PIERER Innovation research team. The team therefore seized the opportunity to participate in the AERO project, which is extending and optimizing the software stack for cloud deployments of EPI's first commercial offering, the Rhea processor commercialized by SiPearl. 'In the future, when the EU has its own, competitively priced cloud processors, even major cloud services providers could add them to their offering,' says Christos. 'As an edge-to-cloud digital-twin development company, it's a no brainer for us to be one of the first adopters and put forward our use cases so that the technology is prepared for them and our software can run optimally out of the box.'

Together, these three EU-funded research projects will allow PIERER Innovation to minimize costs and ensure that their code is optimized for EU processors and cybersecurity compliance, while simultaneously contributing to a homegrown semiconductor ecosystem and expanding the thriving European automotive sector. 'We have products on the market that are already doing extremely well. Now, we're looking ahead to see what the next thing will be,' says Christos.

FURTHER INFORMATION

elegant-h2020.eu

secopera.eu

aero-project.eu

ELEGANT has received funding from the European Union's Horizon2020 research and innovation programme under grant agreement no. 957286. SecOPERA and AERO have received funding from the European Union's Horizon Europe research and innovation programme under grant agreement numbers 101070599 and 101092850 respectively.

With online retail showing no signs of slowing, current forms of last-mile delivery are clogging up streets and contributing to noxious emissions. HiPEAC spoke to Gergely Horváth, product lead at LMAD, to find out how the URBANE project, co-funded by the European Union, is testing HeRo (Helsinki Robot) as one solution to this problem, and learn how this model could be replicated in other locations.

When a HeRo comes along

How Helsinki's robot experiment is shaking up logistics

What are the benefits in using robots for last-mile deliveries?

E-commerce is growing in double digits year-on-year, which fuels more and more deliveries. Logistics companies are facing difficulties in maintaining and growing their delivery workforce due to difficult and stressful work environments, while at the same time offering reasonably priced, flexible and eco-friendly deliveries to the end consumers. Autonomous vehicles and delivery robots offer a solution to all these problems: they are fully electric, thus environmentally friendly, and can work 24/7 to offer cost savings and flexibility to both logistics operators and consumers.

What are the main challenges involved in implementing robots for this purpose in cities?

Firstly, the regulatory environment varies in every European country, and even sometimes in every city – this is one key challenge in a mass rollout of autonomous delivery. Secondly, the maturity of the technology is improving, but real autonomy is not there yet, which means that the human workforce overseeing and supporting the operation of these vehicles is still very much needed.

How is the URBANE project contributing towards logistics solutions? How did the project go about putting the robot delivery use case into practice?

In the URBANE project, there are so called 'living labs', which are all testing different innovations aimed at solving logistics challenges. In the Helsinki Living Lab, the focus is on autonomous vehicles and urban logistics consolidation centres. As part of the URBANE project, the local partners are iteratively testing various use cases to evaluate their feasibility in a real urban environment, from construction-site deliveries and business deliveries to regular e-commerce parcel deliveries. Through the project we get to deploy autonomous delivery vehicles in various parts of Helsinki, and the solution is tested with real customers.

Credit: Forum Virium Helsinki

What were some of the main challenges involved in developing the software platform for this use case? What approach did you take to developing this?

The key offering of the LMAD software platform is that it is vehicle agnostic, meaning that it can integrate any autonomous vehicle and is able to manage and control these vehicles during their delivery mission. At the same time, each use case has its own specificities, which require modifications and adaptations to the platform itself. Delivering construction tools and supplies on demand is quite different from delivering books and clothes ordered online. The main challenge is to keep the platform general enough, but also make it flexible enough to be able to handle different use cases and offer a better service to both logistics companies and their customers.

How would you like to see this technology being rolled out in the future? Do you think it has the potential to be applied to any urban environment?

Autonomous delivery will be an increasing part of the last-mile delivery landscape in the coming years as the technologies behind it mature. Pilots and test deployments are key to showcase how these new vehicles will operate in urban environments, and they also pave the way to public acceptance. There will be more forward-looking cities and countries that will take an active role in shaping the regulation to allow these vehicles on the public roads. We also believe that there will not be a single vehicle type that will respond to all delivery needs, but that logistics companies will need a diverse fleet of such robots. To be able to manage such a fleet, a universal solution is needed that is capable of integrating and managing different types of autonomous vehicles. This is the motivation behind LMAD and the software platform we develop.

URBANE has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement no. 101069782.

URBANE puts last-mile solutions to the test in living labs

Paula Kultanen-Ribas, Forum Virium Helsinki

The purpose of the URBANE project is to test various last-mile solutions in cities in order to reduce traffic congestion and emissions caused by urban logistics distribution. The project is testing and developing new solutions, such as autonomous distribution platforms, and seeking commercially viable operating models for last-mile challenges. Funded by the European Union, the URBANE project involves 41 partners from 12 countries.

The new solutions will be tested in Helsinki (Finland), Thessaloniki (Greece), Valladolid (Spain), Bologna (Italy), Karlsruhe (Germany) and Barcelona (Spain). The pilot projects in Helsinki are being carried out by the shipping and logistics company DB Schenker; LMAD, the developer of the software platform designed for operating and managing autonomous robots; and Soben/TwinswHeel, a developer of autonomous solutions, in cooperation with Forum Virium Helsinki and the City of Helsinki.

During spring and summer 2023, pilots involving HeRo, a robot capable of carrying loads of up to 300 kilograms, were carried out at the construction sites of Kalasatama with the materials company Würth Oy, in collaboration with DB Schenker and LMAD. In these pilots, Würth's customer companies could order supplies such as screws or tools and have them delivered to the site by the robot.

The work is based on previous pilots carried out in various projects, such as autonomous delivery robots which were tested in Helsinki in 2021.

Forum Virium Helsinki is a non-profit innovation company owned by the City of Helsinki that creates smart city, smart mobility and data solutions for the cities, growth for startups and companies (like LMAD) and collaboration opportunities for research institutes.

Q-Free supplies technological solutions for tolling and traffic-management, connected intelligent transport systems and connected vehicles. Headquartered in Trondheim, Norway, Q-Free's global footprint reaches from the shores of Europe, to the Americas, Asia, and Australia. HiPEAC caught up with Q-Free Chief Executive Officer Thale Kuvås Solberg to find out more about how the company's solutions can help reduce congestion and contribute to sustainable mobility.

'Our vision is of a free-flowing, clean and safe mobility future'

With a population spread over a landscape spanning almost 400,000km2, including spectacular mountains, fjords and islands, mobility in Norway has always represented a significant engineering challenge. For many years, tolling has been used to finance major infrastructure projects such as bridges, tunnels and roads, providing an alternative to finance generated by oil and gas extraction. Recently, government policy has also aimed to ensure that transportation is more environmentally friendly: Oslo, for example, aims to have almost no emissions by 2030, with a stipulation that all private cars on Oslo's roads will be emission free in 2030, and that all public transport will be emission free in 2028.

This is the context in which Q-Free – a provider of solutions for electronic tolling, traffic management, and connected intelligent transport systems (C-ITS for short) – operates. Connection and collaboration are both key to the company's strategy and a major plank of the sector's sustainability efforts, according to Q-Free chief executive Thale Kuvås Solberg. 'The mobility industry is no stranger to wicked problems, but it seems that we are now maturing in the way we try to solve them, going from an atomistic to a more holistic approach. Political regulations and related funding are increasingly setting a joint direction, and private players are becoming better at stakeholder dialogue and cross-industrial collaboration. This collaborative approach is crucial for more sustainable mobility,' she says. 'The more we understand about how things are connected, the better we understand our role in the mobility ecosystem and the partnerships we should nurture to achieve sustainable value creation at scale.'

This holistic approach drove Q-Free to move away from a portfolio based on business segments towards structuring its solutions into three programmes based on purpose, explains Thale: Q-Flow, Q-Clean and Q-Safe. 'By bundling our offerings under purposes rather than functions, we invited our stakeholders to navigate these offerings from a different perspective and extended our connection towards networks and companies who shared our vision of a free-flowing, clean and safe mobility future,' she says.

Joined-up approach to sustainability

In terms of the trends shaping the mobility sector, although fully autonomous vehicles may be some way off, Thale explains that increased connection has a part to play in greener mobility, citing a study commissioned by Qualcomm Europe that showed that introducing 20% of connected vehicles on European city roads could reduce greenhouse-gas emissions by up to 18% across the European Union, or up to 24% in countries such as Germany (see 'Further reading', below).

Q-Free is headquartered in Trondheim, Norway

'As in every industry, there is a lot of hype in the mobility world. While drone delivery and fully autonomous cars are still more science fiction than reality, the less radical stages of autonomous driving – together with electrification, connectivity, and shared mobility – are some of the maturing trends that will continue to shape our industry,' she says.

Context is key to how the mobility sector is evolving in different areas, Thale says. 'New and enabling technologies will fuel a constant stream of mobility innovation; the industry is being shaped differently in different markets, depending on the existing infrastructure, demographics, values and behaviours, economy,

Q-Free Kinetic® Mobility platform for traffic management

climate change, and public policies.' This means that a solution that is mature in one market can be groundbreaking in another, she adds. The response of Q-Free, says Thale, is to consolidate its existing technology, investing in the development of a C-ITS platform that connects vehicles and subsystems while enabling applications to be tested in real-life settings, while at the same time using new technologies such as artificial intelligence to future-proof solutions.

Q-Free's main product offerings are tolling and traffic-management systems, as well as active transportation systems. 'We currently have around 40,000 intersection controllers, 50,000 active Intrada automatic licence-plate recognition (ALPR) licences, 2,000 toll lanes, and traffic management for over 500,000 lane miles of highway in operations across the world,' says Thale. 'These intelligent solutions pave the way to more efficient, safe, and environmentally friendly transportation.' The company provides clients with full solutions including hardware, firmware and software.

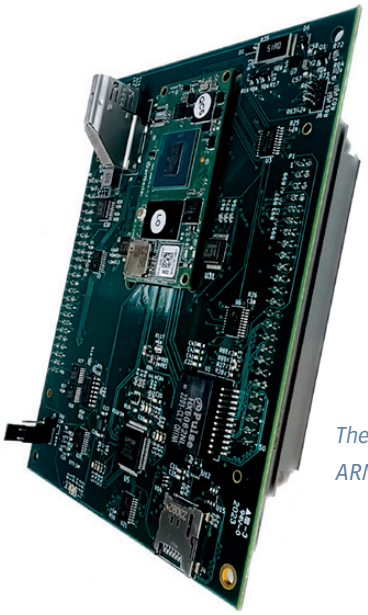

The Q-Free AI-enabled Velocity™ ARM processor

As an example, Thale cites Q-Free tolling systems in Bangkok, one of the three most-polluted cities in the world, which she says have reduced the time it takes to pay a toll from between four and 15 minutes to three seconds. 'With 5.7 million cars travelling through these systems every day, this saves 120,000 hours of idle transportation,' she notes. Meanwhile, the company has been working with the University of Melbourne, Lexus and ambulance services in Victoria to solve issues around connected vehicles for emergency transportation.

However, according to Thale, the company also tries to integrate environmental and social sustainability into its operations, as well as its solutions. 'In the context of supply chains and material flows, our supply management plays a key role in addressing environmental and social risks and opportunities,' she says. 'If we really want to succeed at moving more people and more goods in a more efficient way, using fewer resources, we need to look beyond the obvious and bring the sustainability focus into all our functions.'

FURTHER INFORMATION

Clue Technologies
'Q-Free won ITS Australia Connected and Automated Vehicle Award'
bit.ly/Q-Free_ITS_award

Partani, S., et al. 'Accelerating Safe and Sustainable Transportation: Smart Cars Communicating with Smart Roads'
Study carried out by the University of Kaiserslautern (TUK) and German Centre for Artificial Intelligence (DFKI), sponsored by Qualcomm Europe
bit.ly/Acc_Safe_Sustainable_Trans_2023

Advanced computing is a key part of future air mobility. In this article, Jon González, head of innovation at Clue Technologies, sets out how the company is making advances in this area.

A Clue to the future of air mobility

Comprising a research arm (Clue Technologies) and manufacturing arm (Clue Aerospace), the Spanish company Clue focuses on computing systems for the aerospace sector. Since its establishment in 2013, this small / medium enterprise (SME) has been carrying out research and development (R+D) into onboard computing and high-performance processing architectures, while also ensuring that this work meets the stringent safety requirements of the aerospace industry. According to Jon González, the company's head of innovation, one of Clue's standout features is that it takes a comprehensive approach, covering the entire lifecycle of its products. 'Clue is the sole integral developer of avionics in Spain and one of the few European companies that can cover all development stages in house,' he notes.

Using R+D as a basis to progress its products, the company participates in initiatives such as the European Union's Clean Sky 2 programme, in which it developed modular integrated avionics solutions tailored to small aircraft. Another project, funded by the Spanish government, involved accelerating avionics systems using general-purpose graphics processing units (GPGPUs) for massive data analysis. In addition to R+D, Clue also undertakes industrialization, manufacturing, marketing, and support of equipment based on the technologies developed, Jon explains.

Intelligent applications for air, land and sea are the main areas for the company's equipment. In particular, urban air mobility (UAM) is a primary field of application, says Jon, with Clue currently participating in research projects focusing on software

Future urban air mobility, as shown in this image, is a key area for Clue Technologies © NASA / Lillian Gipson

training and hardware-acceleration architectures for deep-learning-based autonomous systems for UAM. 'Partly due to growing traffic congestion issues in cities, this market has great potential. One study estimates that it could grow from around $2.91 billion in 2023 to around $36.75 billion in 2035,' he notes.

Several research projects in which the company participates cover hardware and software technologies for autonomous flight systems, where Clue aims to define and validate advanced architectures encompassing computing, sensors, and machine-learning applications. This is paving the way for intelligent systems with hybrid navigation capabilities, such as vision navigation.

One example is SKY AI CONNECT, funded by the Spanish Ministry for Innovation and delivered by a consortium comprising Clue Technologies (which leads the technical development of the computing platform), Capgemini Engineering and Rapita Systems SL. The project team also includes Barcelona Supercomputing Center and FADA-CATEC as subcontracted research centres, Clue Aerospace as the demonstrator manufacturer, and Airbus DS as a collaborating company.

'One of the major topics in SKY AI CONNECT is advanced computing and intelligent communication strategies for urban air mobility (UAM), aiming to enhance data mining and automation in both manned and unmanned aircraft using artificial intelligence (AI),' Jon explains. 'Clue's key contribution is in developing a safety-critical (DAL-A) airborne computer that provides ARM-based multicore processing and low-power AI acceleration. This will contribute to improved failure prediction algorithms, increased UAV autonomy, and enhanced communication for mission success in UAM applications.'

"Urban air mobility could grow from $2.91 billion in 2023 to $36.75 billion in 2035"

Clue is also involved in research to develop AI training tools for vehicle autonomy, says Jon. 'The key objective of this project, known internally as Ghost Rider, is to provide extensive synthetic and real-time data to train deep-learning algorithms, which will serve as the foundation for fully autonomous unmanned aerial vehicles (UAVs). Clue's platform solution will assist manufacturers in data collection and AI training, as well as the testing, validation and deployment of fully autonomous systems that will enable the U-space – the airspace, services and systems that allow for mass unmanned aircraft systems (UAS) transit – and UAM of the future,' he adds.

The company's ability to independently manufacture and assemble prototypes resulting from these initiatives is a major benefit, says Jon, and Clue also provides the testing environment for methodologies and artefacts to help achieve certification within standards set by the European Aviation Safety Agency (EASA) and other regulatory bodies.

'In the words of renowned computer researcher Alan Kay, "people who are serious about their software should make their own hardware". Clue embodies this philosophy as we advance into the realm of autonomous flight systems for UAM,' concludes Jon.

FURTHER INFORMATION

Clue Technologies
clue.aero

Clean Aviation Joint Undertaking
clean-aviation.eu

Urban Air Mobility (UAM) Market - Growth, Trends, COVID-19 Impact, and Forecasts (2023 - 2035)

Mordor Intelligence
bit.ly/Mordor_UAM_market_2018-2035

Funded by the European Union's Horizon 2020 research and innovation programme, the AMPERE project sought to bring high-performance parallel computing concepts to cyber-physical systems, applied to rail and automotive use cases. In June, the project celebrated its final event as a webinar hosted by HiPEAC. This article highlights some key takeaways from the project.

Turbocharging mobility with high-performance computing: The AMPERE project

Concluding in June 2023, the three-year AMPERE project developed a software framework which sought to fully exploit computing resources for cyber-physical systems through high-performance parallel programming models.

As the HiPEAC community is well aware, in the high-performance computing (HPC) field researchers have been trying to extract optimal performance from increasingly complex hardware – including multicores and different kinds of processor such as graphics processing units (GPUs) and field-programmable gate arrays (FPGAs) – for some time. As explained by AMPERE coordinator Eduardo Quiñones (Barcelona Supercomputing Center), the project's main challenge was to apply these techniques to cyber-physical systems (CPS). Due to their interaction with the real world, CPS have the added detail of non-functional requirements including real-time responses, keeping within tight energy budgets, ensuring user safety, and robust protection from cybersecurity attacks.

To reduce the so-called 'software productivity gap', where software advances fail to keep pace with the increasing power of hardware resources, HPC uses parallel programming models. These have a number of advantages, including abstraction which hides hardware complexity and the ability to port across architectures. Meanwhile, in embedded computing, the concept of model-driven engineering (MDE) refers to a technique to make systems 'correct by construction': the system is described and the code is automatically generated, allowing formal verification. However, this technique, until recently, was only for single-core execution, or at most with very limited parallel support.

AMPERE set out to bridge the gap between model-driven engineering and parallel programming, with the aim of making advances on several fronts, as follows:

- **Code synthesis methods** which, given the description of the system in a domain-specific modelling language, could efficiently generate parallel source code while maintaining non-functional requirements and composability.
- **Runtime parallel frameworks** that guarantee system correctness and exploit the performance capabilities of parallel architectures.
- **Integration** of parallel frameworks into MDE frameworks.

The main achievements of the AMPERE project were as follows:

- Redefining HPC programming models to unleash their potential for CPS, thus streamlining development and reducing time-to-market for parallel and heterogeneous processor architectures.
- Enabling the seamless transformation of domain-specific modelling languages (DSMLs) into the OpenMP programming model for efficient exploitation of parallel resources.
- Addressing non-functional requirements to ensure the resilience, safety, and security of CPS.
- Providing safety and security mechanisms through a hypervisor and operating systems, PikeOS and OpenERIKA while efficiently supporting the parallel execution model.

Members of the AMPERE project consortium

How model-driven engineering (MDE) enhances productivity

According to Michael Pressler of AMPERE partner Bosch, MDE has a number of benefits, including improving insights into dynamic system behaviour (both systems in development and those deployed in the field), assessing design choices and requirements before implementation, and identifying opportunities in areas such as scheduling parameters, mapping decisions, and so on.

In AMPERE, two domain-specific modelling languages were used:

- **Capella**, used by Thales, is an open-source framework for high-level system modelling in the rail domain.
- **AMALTHEA**, an open-source tool platform focusing on the automotive domain which provides a detailed definition of the physical architecture.

Researchers from Thales Research and Technology also developed a bridge between the two modelling languages to keep them in synch, allowing updates while keeping the models intact.

AMPERE project use cases: automotive and rail

The AMPERE software framework's applications were evaluated through two use cases in the automotive and rail industries. Both cases require swift and accurate responses to increasingly complex inputs. In the automotive use case, the Predictive Cruise Control (PCC) extended adaptive cruise control by calculating future vehicle velocity using electronic horizon data to enhance fuel efficiency. Using the AMALTHEA DSML and optimized OpenMP parallel source code, it effectively exploited NVIDIA Jetson Xavier AGX and Xilinx UltraScale+ platforms while meeting defined non-functional requirements.

The rail use case implemented Thales' Obstacle Detection and Avoidance System (ODAS) to evaluate AMPERE technology using the Capella DSML and the AMPERE bridge for model-to-model transformation. By transforming the model into OpenMP with AMPERE synthesis tools, it harnesses the parallel capabilities of multicore, heterogeneous platforms with GPU and FPGA acceleration, while ensuring non-functional requirements.

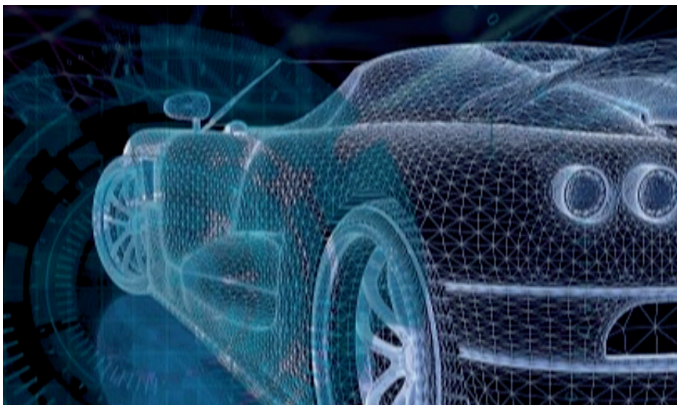

'The AMPERE project provides a leap forward in the cyber-physical computing domain by enabling HPC parallel programming models to offer a faster and more reliable and resource-efficient development and deployment of systems, like rail and automotive. The applications of the AMPERE software framework could potentially revolutionize a wide array of industries that require the use of sophisticated parallel and heterogeneous computing technologies,' commented HiPEAC member Eduardo Quiñones, who leads the Predictable Parallel Computing Group at BSC.

AMPERE: Key achievements

30% reduction in software development costs, while providing the required performance and energy budget imposed by system.

Up to **three times performance speedup** and a **system utilization of 100%** for the two AMPERE use cases, guaranteeing the fulfilment of the non-functional requirements.

Provide **extensions for automotive and railway DSMLs** to better capture requirements.

New extensions to the OpenMP parallel programming framework targeting cyber-physical systems.

The AMPERE webinar is available to view on the HiPEAC TV YouTube channel: bit.ly/AMPERE_final-event_video

AMPERE received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no. 871669.

AMPERE
A Model-driven development framework for highly Parallel and EneRgy-Efficient computation supporting multi-criteria optimisation

What are the main building blocks of tomorrow's intelligent transport systems? In this article, Vladislav Kashansky and Sara Agha Hossein Kashani explain how the data-analytics and computational software developed by Austrian deep-tech company Eteronix contribute to smarter, more sustainable mobility and transport.

HPC, data processing, cloud and edge power: The ingredients of intelligent transport systems at Eteronix

The integration of new information and communication technologies into the field of intelligent transport systems is crucial when working towards sustainable transportation – one of the major fronts in the combat against climate change. To ensure safe and efficient transport, this field relies heavily on the collection, real-time processing, and analysis of data.

Accurate real-world data is essential for a multitude of purposes, including the estimation of existing transport processes' efficiency and further multi-organization management. It is used to establish relevant scenarios, connect simulation results to reality, and set up realistic large-scale transport models. Big-data processing is also critical for estimating delivery times, monitoring sensor data, managing large-scale optimal routes, and reducing carbon emissions to address the challenges posed by the climate change.

In conjunction with data processing, there is a need for efficient green and scalable optimization platforms, which are especially important in the context of large-scale real-world scenarios. Our main objective can be described by the attempt to address several Sustainable Development Goals (SDGs) of the United Nations (UN). It is also important to create simulated environments that accurately represent real-world scenarios and enable developers to train and test algorithms in a safe and controlled environment before deploying them on actual transportation ecosystems.

Achieving the smooth and efficient operation of intelligent transport systems is likely to depend on solving the key issue of worldwide system interoperability between standardized and secure computational platforms. This demands an upgrade of existing cyber-physical infrastructures, with an emphasis on ensuring the future maintenance of their components.

The role of Eteronix in intelligent transport systems

An Austrian deep-tech company, Eteronix targets small / medium enterprises, research institutions, and government organizations, specializing in the technological aspects of digital transformation. Eteronix acts as a smart systems architect and integrator in the domains of high-performance computing (HPC), data analytics, and cloud / edge computing. The company develops analytics and computational software, as well as large-scale data processing systems, capable of functioning in complex industry scenarios.

The Eteronix research laboratory conducts on-demand research-as-a-service (RaaS)-based analysis of complex networks and distributed cyber-physical computing systems. The company also synthesizes and searches for novel forms of adaptive optimization algorithms. In addition, the team designs computational software which is then integrated with advanced artificial intelligent (AI) and HPC technologies, as well as with existing large-scale data processing platforms.

Recently, Eteronix contributed to a research manuscript presented at the ACM GoodIT 2023 conference. In collaboration with colleagues at the University of Calabria, Chinese Academy of Sciences, UC3M and Dtok Labs, the team analysed multi-objective optimization aspects in the logistics sector – an integral part of intelligent transport systems. This work considers transportation costs, emissions, order importance and risks for failures for the logistic route-order dispatching system.

FURTHER READING

'Intelligent and Sustainable Transportation through Multi-Objective Model for the Logistic Route-Order Dispatching System'
GoodIT '23: Proceedings of the 2023 ACM Conference on Information Technology for Social Good – September 2023 – pp. 530–536
doi.org/10.1145/3582515.3609578

As part of our series focusing on the most innovative small / medium enterprises (SMEs) on the European technology scene, this article presents Clue Technologies, a company which combines high-performance computing with artificial intelligence (AI) for safety-critical industries.

Safety and performance for next-generation avionics: Clue Technologies

COMPANY: Clue Technologies S.L.
MAIN BUSINESS: Research, development and manufacture of high-performance computers for safety-critical industries
KEYWORDS: high-performance computing, smart avionics systems, airborne artificial intelligence, real-time data analytics
LOCATION: Málaga, Spain
Website: clue.aero

Founded in 2013, the Spanish SME Clue is well known for its innovative embedded computing systems for the aerospace sector. Over the past decade, Clue has undertaken significant research and development (R+D) into high-performance computing (HPC) platforms, focusing on safety and massively parallel processing for critical applications while simultaneously developing AI-based software that harnesses the full potential of this infrastructure.

With a diverse range of expertise in HPC, airborne AI, real-time data analytics, and safety-critical systems, Clue is continuously engaged in pioneering research on comprehensive hardware and software platforms tailored to the aerospace industry.

One of the key strengths of Clue lies in its development of embedded applications focused on 'airborne artificial intelligence' (A^2I), specifically designed for seamless integration into onboard systems. The Clue WittyBox™ family of flight and mission computers serves as a robust engine for data analytics applications, efficiently managing massive data capture across diverse networks. By leveraging these tools, new flight and mission applications can analyse complex problems in real time and make automatic decisions.

The Clue WittyBox Calibre range is a compact, robust form factor

In the realm of innovative mobility projects, as explored on pp.18-19, Clue has also dedicated significant R+D efforts to the creation of an integrated modular avionic computer for small aircraft and unmanned aerial vehicles (UAVs) within the Horizon Europe and Clean Sky 2 frameworks. The company is currently focused on the urban air mobility (UAM) field, actively working towards low-power, mass-parallel computing systems to enable intelligent capabilities while using a design methodology that enables the company to provide safety evidence and artefacts to bolster confidence in the certification potential of these technologies with the European Aviation Safety Agency (EASA).

Another notable project in this field is the Clue platform solution for the development of autonomous systems of systems (SoS). This platform facilitates dataset generation, deep reinforcement, testing, validation, and deployment of fully autonomous systems, combining AI and sensor fusion to process information and enable autonomous decision-making capabilities.

This open platform will allow developers to leverage a full software stack to create their own mobile autonomy applications. It will also be scalable, from a single system to a complete system of systems, and will meet functional safety standards, including aerospace standards, future EASA requirements for UAM, and trusted AI best practices.

Through these groundbreaking initiatives, Clue aims to become the premier European supplier of reliable training for autonomous navigation UAM systems, solidifying its position as one of the few companies worldwide with this specialization. With the ability to integrate AI and big-data-based systems highly sought after by the U-space industry and the UAM sector, Clue is poised to embark on numerous transformative projects in the next future.

Innovation Europe

In this issue's focus on the latest publicly funded computing systems research, we learn about smart cloud scheduling, operating systems for an interoperable compute continuum, and a new internet-of-things architecture. Plus a new centre of excellence focusing on exascale computational fluid dynamics and a project bringing model-based design to open hardware in space.

SMART SCHEDULING FOR COMPLEX CLOUDS: THE DECICE PROJECT

Cloud computing has grown enormously over the last decade, transforming the way we compute and driving powerful new applications such as smart cities, industrial automation and data analytics. However, this expansion has been accompanied by increased complexity, for example:

- a more varied hardware landscape, with different platforms specialized for different uses;
- requirements relating to embedded and cyber-physical computing, including location awareness, ultra-low latency, and security;
- increasingly complicated cloud application scenarios running on distributed systems, divided into components that must run with high availability.

Unifying such diverse systems into centrally controlled compute clusters and providing sophisticated scheduling decisions across them are two major challenges in this field. Scheduling decisions for a cluster consisting of cloud and edge nodes must consider unique characteristics such as variability in node and network capacity. A common solution for orchestrating large clusters is Kubernetes. However, Kubernetes is designed for homogeneous clusters and, while there are many applications and extensions available, none accounts for the optimization of both performance and energy or addresses data and job locality.

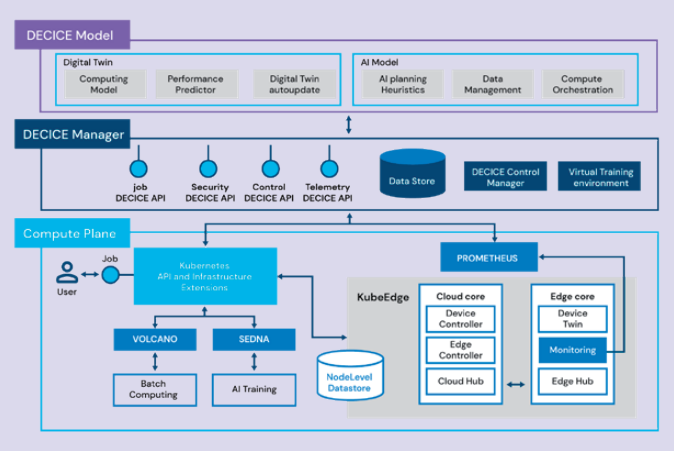

Bringing together 13 partners from Austria, Germany, Italy, Sweden, Turkey and the United Kingdom, the European Union-funded DECICE aims to remedy this situation. The project is developing an open cloud-management framework that optimizes applications automatically by mapping jobs to the most suitable hardware resources in a heterogeneous system landscape spanning edge, cloud and high-performance computing (HPC). This is done using an artificial intelligence (AI) scheduler that makes decisions on the placement of jobs and data, as well as rescheduling jobs in response to system changes.

In parallel, using holistic monitoring, the project is also constructing a digital twin of the system to create a virtual training environment. This training environment generates data for the training of machine-learning models and the exploration of 'what if?' scenarios.

As part of the project, the portable framework is being integrated into the Kubernetes ecosystem and validated using real-world use cases in fields including smart cities, medical imaging and disaster response on real-world heterogeneous systems.

PROJECT NAME: DECICE: Device-Edge-Cloud Intelligent Collaboration framEwork
START/END DATE: 01/12/2022 – 30/11/2025
KEY THEMES: cloud computing, edge computing, artificial intelligence
PARTNERS: Germany: Georg-August-Universität Göttingen (coordinator), Gesellschaft für Wissenschaftliche Datenverarbeitung MBH Göttingen (GWDG), University of Stuttgart, Huawei Technologies Düsseldorf GmbH; Austria: Forschung Burgenland GmbH, SYNYO GmbH; Sweden: Kungliga Tekniska Hoegskolan; Italy: E4 Computer Engineering SPA, Consorzio TOP-IX, Universita di Bologna; Turkey: Marmara University, BigTRI Bilişim; United Kingdom: The Numerical Algorithms Group Limited
BUDGET: € 5.6 million
Website: decice.eu Twitter: @DECICE_EU
LinkedIn: linkedin.com/in/decice-project-b0b55b25a

DECICE has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement no. 101092582.

ASSIST-IOT: NEXT-GENERATION INTERNET-OF-THINGS ARCHITECTURE ENABLES SAFER, SMARTER REAL-WORLD USE CASES

ASSIST-IoT is a European-funded research and innovation action that aims at creating a novel next-generation internet-of-things (NG-IoT) architecture addressing scalability and flexibility of data processing and analytics. To do this, ASSIST-IoT is implementing a decentralized and secure solution, supported by smart components (i.e. devices, edge nodes, networks, applications, and services), allowing the execution of context-aware applications with new interaction interfaces, such as augmented reality (AR) / virtual reality (VR) and mixed reality (MR). The blueprint architecture to be delivered by the project covers the emerging requirements of NG-IoT deployments.

Members of the ASSIST-IoT consortium at Malta Freeport Terminal

In addition to establishing the basic technology for such solutions, ASSIST-IoT will indicate how to incorporate new 'modules' (enablers) to power various NGIoT innovations. The solution integrates artificial intelligence (AI)-based functions transferring intelligence closer to the edge (data sources), including devices. The reference architecture is being validated in three realistic pilots: (i) port automation; (ii) smart safety of workers, and (iii) cohesive vehicle monitoring and diagnostics. Each pilot includes different scenarios, in which different technological pillars and enablers are executed and validated. Experiences from pilots will be used to improve action outcomes (feedback-loop), guaranteeing quality and a broad range of applicability for the results.

The breakdown of the pilots is as follows:

- **Pilot 1: Port automation**
  - Tracking assets in the terminal yard
  - Automated cargo handling equipment cooperation
  - Crane remote control with AR support
- **Pilot 2: Smart safety of workers**
  - Occupational safety and health monitoring
  - Fall arrest monitoring
  - Safe navigation
  - Health and safety inspection support
- **Pilot 3: Cohesive vehicle monitoring and diagnostics**
  - Advanced powertrain monitoring and diagnostics
  - Vehicle condition monitoring

One of the main objectives is to build this architecture following a human-centric approach. This means, in the context of IoT, that the humans using / being part of the IoT deployment are at its very centre, in terms of usability, data protection, rights and freedom preservation, and the computing itself. To ensure this human-centred approach, ASSIST-IoT incorporates a suite of technologies, including distributed ledger technology (DLT) for decentralized cybersecurity and privacy, as well as distributed AI enablers, into its architecture. The project is also working on human-centric tools and interfaces.

PROJECT NAME: ASSIST-IoT: Architecture for Scalable, Self-*, human-centric, Intelligent, Secure, and Tactile next generation IoT
START/END DATE: 01/11/2020 – 31/03/2024
KEY THEMES: artificial intelligence, computing continuum, internet of things
PARTNERS: Spain: Universitat Politècnica de València (coordinator), Prodevelop, S21Sec; Poland: Systems Research Institute (Polish Academy of Sciences), Central Institute for Labour Protection – National Research Institute, Mostostal Warszawa, Orange Polska; Greece: Centre for Research & Technology, Hellas – CERTH, INFOLYSiS, Institute of Communication and Computer Systems (National Technical University of Athens); France: Terminal Link; Netherlands: Neways Technologies; Finland: Konecranes Finland; Germany: Ford-Werke, Twotronic
BUDGET: €7.9 million
Website: assist-iot.eu
Twitter: https://twitter.com/@AssistIot
Facebook: facebook.com/assistiot
Instagram: instagram.com/assistiot
LinkedIn: bit.ly/ASSIST-IoT_LinkedIn

A UNIFIED OPERATING SYSTEM FOR THE INTERNET OF THINGS: AEROS

One of the main messages of the HiPEAC Vision is the need for orchestration to overcome fragmentation and provide a more holistic, enriching experience of the future web, including the internet of things (IoT). Here, Harilaos Koumaras and Vassilis Pitsilis (Institute of Informatics and Telecommunications, National Centre of Scientific Research 'Demokritos'), explain how aerOS is developing a meta operating system (meta-OS) to efficiently orchestrate compute and network resources.

Anyone surveying the current internet of things (IoT) landscape would be struck by how fragmented it is. 'What we have currently is a plethora of isolated processing units and private computing islands. These lack the necessary resources and services for holistic solutions in the IoT, either in an industrial or personal setting,' explains Harilaos Koumaras. 'Numerous computing units and private networks act as data concentrators and forwarding equipment to store and process a large volume of data in central, commercial cloud infrastructures operated by a small set of service providers. This deprives vertical IoT stakeholders from having full control of their services and governance of their data.'

Another problem is that, although many services already exist, these cannot be reused for most industry verticals, leaving companies to develop solutions from scratch to solve similar problems, adds Vassilis Pitsilis. 'There is no lingua franca for IoT service developers, which would allow a common, underlying layer where existing solutions can serve similar needs, meaning that companies are forced to fully adapt their existing operating runtime environment.'

To address these issues, the aerOS project is developing a meta-OS which will unify and orchestrate compute and network resources in the most efficient way. Delivered by a consortium of 27 partners from 11 countries, the project will provide a unified execution environment for IoT service developers across a distributed computing environment. 'The idea is that end users – that is, IoT developers from a variety of industry verticals – should have the impression of seamless integration of these underlying resources,' says Harilaos. 'Users should be able to take advantage of smart aerOS orchestration and federation decisions, deploying services as close to the edge as possible while leveraging the most appropriate and efficient combination of resources across the continuum, from edge to cloud.'

How it works

Diverse compute and network resources are integrated into what the project names 'infrastructure elements', the building blocks of the aerOS continuum. In turn, these infrastructure elements form aerOS 'domains', as Vassilis explains: 'There are two prerequisites that make a set of connected compute and network resources an aerOS domain. The first is that these compute resources are integrated as infrastructure elements, which we refer to as IEs – that is, that they are able to support workload execution and provide a minimum set of aerOS integration capabilities, such as manageable network functionality. The second is that a set of cores aerOS services is deployed on top of these IEs.'

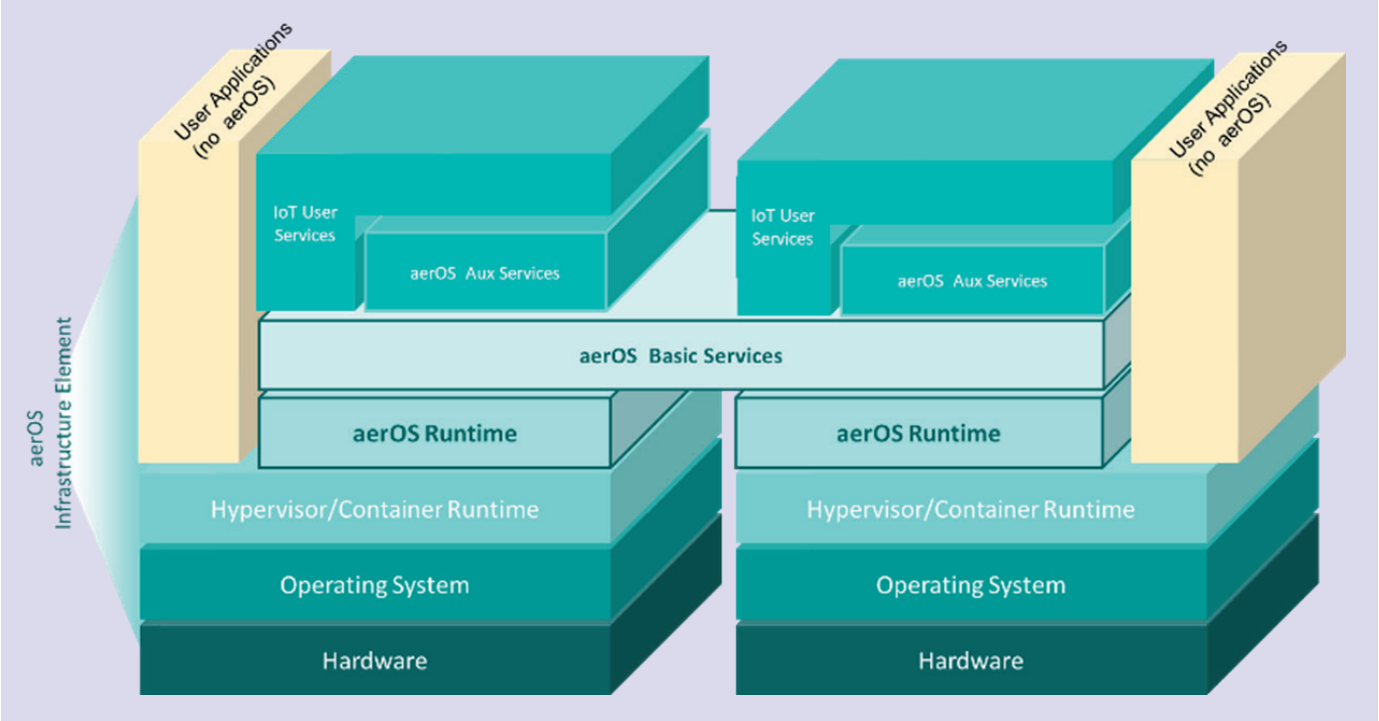

These core aerOS services can be broken down into four categories, Harilaos adds, as follows:

- **Federation service**, provided by the federator component. 'This enables state propagation of domains and IEs across the continuum, as well as forwarding orchestration requests,' says Vassilis. 'Federator requests and information can be propagated between domains thanks to the mechanism implementations employed (NGSI-LD).
- **Orchestration service**: This service receives descriptions of users' intentions regarding IoT service deployment. It encompasses decision-support systems and trust-management services which are translated to actual deployment requests, and finally deployments on IEs. This is a two-level process: a high level receiving the user request and handling all the decision complexity, and a low level to access resources,' explains Vassilis. 'It is important to note that, although decision-support systems are deployed locally – i.e. within the specific domain – they leverage information regarding the state of domains and IEs across the whole continuum. This information is provided "for free" by the data-fabric services, as explained below.'
- **Data-fabric services**: implemented within each aerOS domain, these manage the identification and interoperable integration of the data, before enforcing the required governance policies, says Vassilis. 'As data fabric is a layer running across the continuum, they also provide all the mechanisms to discover, connect and retrieve data from other aerOS domains. Consumers should therefore be able to ask for data and get them in a transparent way without knowing how or where they came from.' The term 'consumers' in this sense could refer to other aerOS services (such as the orchestrator), IoT applications, or even external agents (such as the portal),
- **Security services**, which integrate authentication, authorization and access policies based on roles and identities.

In addition to these core services, additional services enable aerOS to act as an intelligent and secure meta-OS able to manage the continuum, such as:

- **Trust-management services**: These can exploit the aerOS data-fabric provisions and calculate a trust score for each aerOS domain – or even node – to guide the domain orchestrator to the most appropriate choices.
- **AI decision-support services**, which can leverage data retrieved by the data fabric to provide input to the orchestrator regarding the best placement (either locally or in another domain).
- **Analytics services**, which can provide insights into data, decision-making, and data processing for other components on request.

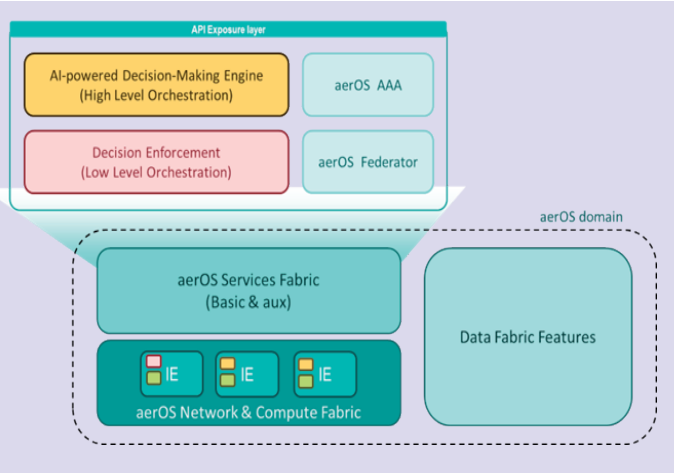

'On top of these services, each aerOS domain exposes a comprehensive, efficient application programming interface (API) to communicate with stakeholders, agents and other domains,' adds Harilaos. The figure above shows a schematic representation of an aerOS domain, with the immediately identifiable continuum enablers highlighted.

Get involved

aerOS has launched a call for proposals for small-scale projects to validate the aerOS architecture, focusing on one of the aerOS use cases. The deadline for sending a proposal is 31 January 2024. For further information, visit the aerOS website: aeros-project.eu/open-calls/open-call-1

PROJECT NAME: aerOS: Autonomous, scalablE, tRustworthy, intelligent European meta Operating System for the IoT edge-cloud continuum
START/END DATE: 01/09/2022 – 31/08/2025
KEY THEMES: operating systems, internet of things, computing continuum, artificial intelligence
PARTNERS: Spain: Universitat Politècnica de València (coordinator), Innovalia, Telefonica, Prodevelop, S21Sec, Nasertic - Navarra de Servicios y Tecnologías; Greece: National Center for Scientific Research 'Demokritos', Cosmote, Fogus Innovations & Services, INFOLYSiS; Austria: TTControl; Germany: Siemens, FIWARE Foundation, John Deere; Cyprus: Eight Bells, Eurogate Container Terminal Limassol, Cyprus University of Technology; Romania: Inqbit Innovations; Ireland: Ericsson; Poland: Systems Research Institute (Polish Academy of Sciences), CloudFerro, Electrum; Finland: ICTFicial; Italy: DS Tech, Università Politecnica di Milano, Made; Switzerland: Switzerland Innovation Park Biel / Bienne
BUDGET: € 11.8 million
website: aeros-project.eu
Twitter: @AerosProject
Facebook: facebook.com/aerosproject
LinkedIn: linkedin.com/in/aeros-project

aerOS has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement no. 101069732.

FINDING SOLUTIONS TO GRAND CHALLENGES AT THE FRONTIER OF CFD

The Center of Excellence in Exascale CFD (CEEC) project, which is co-funded by participating countries and the EuroHPC Joint Undertaking, launched in January 2023. Its objectives are to scale existing European flagship computational fluid dynamics (CFD) codes to run simulations that require an entire exascale HPC system like LUMI in Finland and those being procured by the EuroHPC Joint Undertaking.

Building on the work of prior and complementary European projects, CEEC will use both established and novel techniques and algorithms to push five CFD codes to exascale performance. The codes are: Alya, FLEXI, Nek 5000/Nek RS, NEKO, and waLBerla. To meet the desired performance objectives, CEEC will improve the efficiency of accelerator exploitation, adaptive mixed-precision algorithms, uncertainty quantification, and data compression at scale. It will also implement hybrid remote visualization, large-scale topology optimization algorithms, and will improve the prediction of turbulent flows at high Reynolds numbers (which are used to identify fluid behavioural patterns) via machine-learning-based sub-models.

With these code improvements, CEEC will run six lighthouse use cases with high impact on industrial applications including increasing aeroplane safety and efficiency, creating the first public high-performance computing (HPC) models of static mixers, increasing the success of off-shore wind farm installation, modelling atmospheric boundary layer flows, and predicting the turbulent flow field around large merchant ship hulls. Alongside this work, CEEC will also offer a range of community workshops, webinars, and presentations or mini symposia, both self-hosted and at familiar HPC events, over the next four years to transfer project knowledge to the broader CFD community.

Future updates and events will be available on the project website and social media channels – see details below.

> ***"CEEC will improve the efficiency of accelerator exploitation, adaptive mixed-precision algorithms, uncertainty quantification, and data compression at scale"***

Members of the CEEC consortium

PROJECT NAME: CEEC: Center of Excellence in Exascale CFD
START/END DATE: 01/01/2023 –31/12/2026
KEY THEMES: big data, analytics and management; energy efficiency; exascale computing; optimization and scalability; programming models and tools; visualization; machine learning; aeronautics; wind energy
COORDINATOR: Kungliga Tekniska Hoegskolan (KTH)
PARTNERS: Aristotelio Panepistimio Thessalonikis (AUTH), Barcelona Supercomputing Center-Centro Nacional de Supercomputacion (BSC), Friedrich-Alexander-Universitaet Erlangen-Nuernberg (FAU), University of Stuttgart (USTUTT), Umea Universitet(UMU), Danmarks Tekniske Universitet (DTU), Bundesanstalt fuer Materialforschung und -pruefung (BAM)
BUDGET: approximately €7.8 million
Website: ceec-coe.eu
Twitter: twitter.com/CEEC_CoE
LinkedIn: bit.ly/CEEC_LinkedIn

Funded by the European Union. This work has received funding from the European High Performance Computing Joint Undertaking (JU) and Sweden, Germany, Spain, Greece, and Denmark under grant agreement no. 101093393. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European High Performance Computing Joint Undertaking (JU) and Sweden, Germany, Spain, Greece, and Denmark. Neither the European Union nor the granting authority can be held responsible for them.

METASAT: PIONEERING MODEL-BASED DESIGN FOR OPEN-ARCHITECTURE HARDWARE IN SPACE AND AVIATION

Launched in January 2023, the EU-funded METASAT project will provide a holistic and modular model-based framework for designing and testing software modules targeting open, high-performance hardware in the space and aviation domain.

Bringing together the expertise of knowledge generators such as Barcelona Supercomputing Center (the project coordinator), Ikerlan, and Collins Aerospace, along with technology integrators fentISS, and OHB, one of the main end users from the space sector, the METASAT project is set to revolutionize the model-based design toolchain landscape. By introducing innovative software virtualization layers leveraged by XtratuM hypervisor, the consortium aims to optimize the development and deployment of final applications.

The impetus behind METASAT is the pressing need to address the growing complexity of new satellite designs and the demand for reliable on-board software technology. To tackle this challenge, the consortium has been developing a new model-based engineering methodology for designing software modules.

At the core of METASAT's vision is the creation of a toolchain that enables the design of software modules using high-level modelling languages, tailored for the specific hardware / software layer. By doing so, the project aims to make a substantial impact on the time and cost of developing new systems, mitigating the rising costs driven by system complexity. This reduction in costs not only enhances competitiveness and innovation within the industry but also reinforces dependability and reliability across the board.

Multifaceted approach to satellite software design

METASAT's partners contribute in different ways to the model-based engineering methodology for satellite software design. Barcelona Supercomputing Center (BSC), the project coordinator,

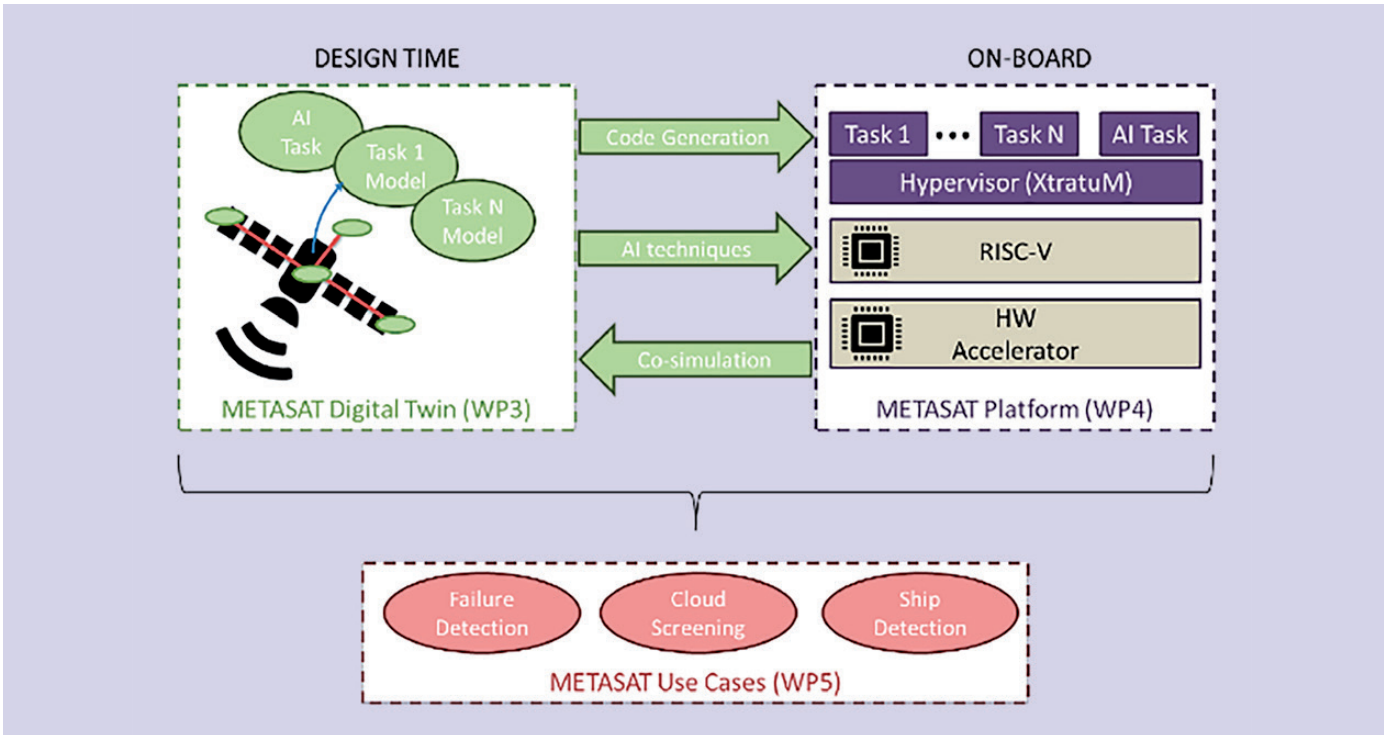

provides a prototype, high-performance hardware platform, based on RISC-V, for future on-board processing. This prototype is used to demonstrate project developments on a virtual platform and a field-programmable gate array (FPGA) implementation. The architecture is a multicore based on Frontgrade Gaisler's NOEL-V processor, enhanced with the SPARROW artificial intelligence (AI) single instruction, multiple data (SIMD) unit – itself developed with the support of European funding – and coupled with the RISC-V based Vortex graphics processing unit (GPU). In addition to hardware design, BSC brings expertise in the use of GPUs in certified safety-critical environments and two open-source use cases for on-board machine learning, based on the European Space Agency (ESA) OBPMark-ML benchmarking suite.

Meanwhile, fentISS is supporting the hardware platform in the XtratuM hypervisor and the RTEMS board support package (BSP) for XtratuM. Additionally, both bare-metal (XRE) and RTEMS partitions will be able to access to a network by means of virtual Ethernet devices and an input / output (I / O) server partition. fentISS will also create a GPU virtualization solution to perform AI operations from several XRE partitions.

Collins Aerospace contributes to the development of the model-based engineering (MBE) workflow. In addition, in collaboration with Ikerlan, the aerospace specialists are working on the METASAT Virtual Integration Framework based on open standards (e.g., FMI, TLM 2.0, ED-247) to support the design and the verification of the METASAT digital twins. Collins Aerospace is also investigating AI-based approaches to aid design and testing. For its part, the Advanced Laboratory in Embedded Systems (ALES) brings expertise in MBE methods for design and verification of embedded systems, hybrid systems modelling and analysis, virtual integration methodologies for platform-based design, and the design of industrial standards for co-simulation and model exchange.

"METASAT project will provide a model-based framework for designing and testing software modules targeting open, high-performance hardware in the space and aviation domain"

In collaboration with Collins Aerospace, Ikerlan has started the definition of a toolchain for space applications, following a model-based design (MBD) cycle from requirements to automatic code generation, as well as each step of verification, validation and testing following sector standards. Ikerlan has also started enhancing a state-of-the-art MBD toolchain and is progressing towards the definition of a proof of concept to showcase the toolchain's capabilities. Furthermore, Ikerlan has taken the first steps towards creating the digital twin of space applications.

Finally, OHB, as an end user from the space sector, will be able to test the new design toolchain that will enable the runtime deployment of software modules, in a representative scenario.

Through this project, the METASAT consortium anticipates groundbreaking advancements and contributions to the space and aviation industries. By leveraging cutting-edge technologies and collaborative efforts, METASAT is set to redefine model-based design methodologies and work towards a new era of efficient and reliable software solutions.

PROJECT NAME: METASAT: Modular model-based design and testing for applications in satellites
START/END DATE: 01/01/2023 – 31/12/2024
KEY THEMES: open hardware, RISC-V, hardware acceleration, space applications, model-based engineering
PARTNERS: Spain: Barcelona Supercomputing Center (BSC) (coordinator), fentISS, Ikerlan; Germany: OHB System; Italy: Advanced Laboratory on Embedded Systems (ALES)
BUDGET: € 1,999, 501
CONTACT: Leonidas Kosmidis, coordinator (BSC)
leonidas.kosmidis@bsc.es
metasat-project.eu

METASAT has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement no. 101082622.

For many years, HiPEAC has supported computing-systems researchers in their careers. This tradition continues in HiPEAC7 with dedicated competence-building activities, clustered into two main areas: the ACACES summer school and HiPEAC Jobs careers activities, which we explore in this article.

Build your computing research career with HiPEAC

Since its inception, HiPEAC has striven to be a central hub for researchers over the course of their careers. From funded research stays to internships, from in-person training to online webinars, HiPEAC has tried to provide practical and timely support, whatever your career level or computing systems subfield.

In HiPEAC7, these activities are clustered into two main areas:

ACACES

As reported on pp. 5-6, the annual summer school on Advanced Computer Architecture and Compilation for High-performance Embedded Systems is a treasure trove of expertise in computing-systems knowledge, from safety-critical real-time embedded systems to exascale computing. Spanning a week in July, it is also an excellent opportunity for early career researchers to network with world-class professors and industry representatives in a relaxed, informal setting. The summer school programme is based on the trends and recommendations that appear in the HiPEAC Vision, helping next-generation researchers get up to speed on key themes for the future.

HiPEAC Jobs careers centre

The online HiPEAC Jobs portal showcases a wealth of opportunities in computing systems – on average five per workday. These include permanent paid positions, PhD opportunities and internships. The offers cover all career levels and span the full spectrum from high-performance computing to edge computing, in both academic and industrial settings. They are searchable by key words – try typing in 'machine learning', for example – or you can browse using the filters.

In addition to the digital portal, HiPEAC takes the jobs wall into the real world with the HiPEAC Jobs wall at events – both HiPEAC events such as the conference and summer school, and partner events.

To supplement the jobs portal, HiPEAC also organizes careers-focused events. These include:

- **STEM Student Day** at the HiPEAC conference, in which students may attend the HiPEAC conference for free, meet the cutting-edge deep-tech companies in the exhibition, and network with researchers from all over the world.
- **HiPEAC Student Challenge**, now part of the STEM Student Day: aimed at giving students a taste of what a research career is like, the HiPEAC Student Challenge is an excellent opportunity for participants to develop their programming, research and presentation skills.
- **Careers roundtables**, usually organized during ACACES, where a panel of experts share their personal experiences and career journeys.
- **Inspiring futures sessions**, where researchers can learn about career paths at major companies, research institutions and universities.

Finally, HiPEAC also organizes regular webinars. In addition to dissemination opportunities for researchers, these also aim to help researchers enhance their knowledge and skills, from preparing an impact-driven project proposal to technical skills that are imperative for their day-to-day work. Webinar videos are subsequently shared on the HiPEAC TV YouTube channel, where you can also find videos of ACACES lectures and other enriching technical talks.

Which careers topics are you most interested in? What kind of activities would you like HiPEAC Jobs to offer? Which topics would you like HiPEAC to cover in webinars and other training activities? Let us know – email info@hipeac.net with your ideas.

What drives researchers to undertake an academic career in computer science? Here Panagiota Fatourou, a professor in the Department of Computer Science of the University of Crete, Greece and the Institute of Computer Science (ICS) of the Foundation for Research and Technology - Hellas (FORTH), shares her experience.

Career talk: Panagiota Fatourou

What first sparked your interest in computer science?

When I was finishing high school, computer science was emerging as the science of the future. The more I came into touch with it, the more I realized that it would end up being one of the most influential sciences. Computer science has the power not only to determine future technology, but also to tremendously influence the evolution of many other fields, hopefully with a positive impact on the world. That was something that fascinated me.

When I started my undergraduate studies in computer science, computing technology was experiencing such a dynamic and rapid evolution that it was hard to follow all the technological advances. However, this also meant that there were countless opportunities for progression and learning. Amazingly, more than 25 years later, this is still true.

How did you come to specialize in distributed computing and in algorithms?

Parallel and distributed computing has a long history, dating back to the dawn of operating systems and networks. Although it has been studied for about the same time as computer science itself, the field is rapidly evolving and comes in many flavours, each with its own challenges to consider.

Some of the most pressing challenges are the performance and scalability issues that arise from the increased computing needs of modern applications. To meet these needs, we must exploit the full computational power of modern parallel and distributed platforms, whether in high-performance computing (HPC), in the cloud or at the edge.

This entails harnessing the power of dynamic, large-scale, heterogenous systems, for example by utilizing all their computing elements – not only their multiple nodes, but also the multicore capabilities of each node, as well as the full capacity of the accelerators and devices attached. To do this, we have to utilize powerful mathematical methods and design novel algorithms process large-scale, dynamic computations in a highly scalable and performance- or energy-efficient way. In the process, we have to find answers to fundamental questions that scientists in this field have been struggling to answer for decades.

By specializing in parallel and distributed computing, I had the chance to contribute to an exciting research field that has had significant impact on the evolution of computing in the past, and the potential for even greater impact in the future. My specialization in algorithms reflects my early passion for mathematics and for understanding the properties and limitations of the research problems I study. It also reflects my view that designing efficient algorithms will always be a challenging, intellectually demanding and intriguing task, thus rewarding to those who manage to tame it.

What should associations and projects like HiPEAC do to support computer scientists' careers?

Over the last few years, I have had plenty of opportunities to think about and work on these issues in the context of the Association for Computing Machinery (ACM), where I served as chair of the ACM Europe Council, as well as the founding director of the ACM Europe Research Visibility Working Group (RAISE). Our mission was to increase the visibility of European research and ensure a higher degree of recognition for the achievements of European researchers.

To do this, we proposed a series of actions:

(i) highlight the role of Europe as a hub of research (for example by supporting the organization of more international events, including conferences);
(ii) increase the number of European nominations for awards and distinctions (where these are genuinely deserved);
(iii) increase awareness of the opportunities for recognition of the professional achievements of researchers (for example by identifying stakeholders, such as nominators and nominees, and educating European scholars about the opportunities available and how to take advantage of these);
(iv) ensure geographical coverage in all events, committees, bodies, and activities of the association;
(v) disseminate the achievements of European researchers (for example by publishing interviews, introducing awards, etc.).

The ACM Europe Council has adopted the above directions as part of its official goals and strategy, and is implementing actions along these directions. Even if only partially implemented, this action plan could have a significant impact on the career prospects of European computer scientists and truly boost the visibility of European research.

HiPEAC could play an important role here, for example by raising awareness about opportunities for researchers in Europe (awards, grants, etc.). This could be done by organizing appropriate dissemination activities, but also by instilling a culture among members of the community which is conducive to taking up this kind of opportunities. HiPEAC could also organize and promote targeted events; while the HiPEAC conference is a great initiative in this direction, there is room for other kinds of events. For instance, events bringing leading HiPEAC researchers together with researchers from other disciplines, to identify common ground for collaboration, and events specifically aimed at early career researchers.

Do you have any tips for researchers who are just starting out in their careers, in particular women?

Studying computer science provides excellent employment opportunities, a situation which is likely to continue given the predicted labour shortages in the future. Today, computing is woven into almost all of our activities, so it is a rewarding field of study.

In any case, I think it is important to start your career with passion. Believe in yourself and focus on your professional goals. Be courageous and fight for what you have dreamt of. Don't waste your energy trying to become what others expect you to be. Instead, focus on what inspired you to become a computer

scientist. Be open to knowledge and to different professional experiences; be active and energetic; accept constructive criticism and learn how to push forward.

From a technical perspective, ensure that you have the necessary theoretical and algorithmic background to solve new, challenging problems, and that you know how to express yourself concisely. These are skills that have the power to pave the way to greater success.

Women in computer science work in environments that are heavily male dominated. Like others who are a minority in their field, they may experience professional bias. All kinds of people, irrespective of gender, race, age, disability, sexual orientation, etc., are equally needed in science and should have the same opportunities for professional success. We should get informed about societal stereotypes, which often follow people from childhood, and try to recognize and eliminate them.

The Department of Computer Science at the University of Crete

This year's ACACES included an inspirational careers evening, with a panel featuring the tutors Reetuparna Das, Yulia Sandamirskaya, Axel Nackaerts and Filippo Mantovani, as well as computing student Blanca Sabater. Prior to the panel session, Blanca gave a talk on diversity issues, focusing on neurodiversity and gender. HiPEAC caught up with Blanca to go into more depth.

'We should never assume the other party's mental processes are like ours'

Why is it important to take diversity into account in computing (or any) careers?

I have observed that many jobs involve the execution of a tiny fragment of a huge, abstract project, making us oblivious of the whole and leading to people specializing more and more in less and less. When individuals start to 'classify' themselves in something very concrete, this creates a culture of like-minded people that sets a false standard of what people within a particular branch are expected to be like. These expectations, I believe, are what lead many diverse people, unconsciously, not to take interest in fields lacking relatable role models.

Ultimately, the result of this lack of diversity within workgroups is a constraint on their collective knowledge, potentially causing them to lose sight of the bigger picture, which, in turn, can lead to problems such as bias in artificial intelligence (AI).

In your experience, what sort of problems arise when people aren't aware of diversity issues?

I feel like the main problem in any relationship or interaction between two people is that we assume that the other party processes things mentally in the same way as we do. This is obvious for anyone who has moved abroad, because social constructs change and all of a sudden others don't follow our expectations and we can't make sense of their behaviour.

Autistic people are more likely to digest social constructs in their own unique way, so both parties in an allistic (non-autistic) – autistic relationship / interaction are practically guaranteed to behave in unexpected ways to each other.

I think the best solution to this issue is to create a habit of never assuming that the other party's mental processes are like ours, and making the effort to understand others' mental states by asking them directly.

What do you think people should be sensitive to?

After the two parties have made an effort to understand each other, and recognized that their mental processes differ, it is crucial that neither tries to impose their perspective / behaviour or treat it as a universal truth. Otherwise, it could lead to mistreatment of the other party.

As an example: an autistic person is in the office and can't block the sound of their allistic colleague typing. The autistic person feels overwhelmed, their feelings amplify, they can't focus or disguise their discomfort, and finally they leave the room. In this situation, it would be abusive to condemn the autistic party for their behaviour, since it is not under their control. However, it would be equally abusive if the autistic person tried to modify the behaviour of the other party (asking not to type so loudly, for example) for their own convenience instead of looking for resources and managing their discomfort on their own (by using sound-cancelling muffs, for example).

Autistic people may be unaware of their tone or not understand tone at all, and this is especially true when they feel overwhelmed. In contrast, allistic individuals heavily rely on tone, often altering word meanings and generating behaviour like sarcasm. Autistic people won't understand such behaviour because we prioritize words over tone. Allistic individuals prioritize tone over words and may feel disheartened if the autistic person speaks to them in a more direct manner.

FURTHER READING:

'Neurodiversity as a Competitive Advantage'
Harvard Business Review, May-June 2017
bit.ly/HBR_neurodiversity_competitive

In the latest in our series on PhDs, Lucie Klus tells us about her thesis focusing on energy efficiency and indoor localization for internet-of-things (IoT) and wearable devices, undertaken at Tampere University and University Jaume I.

Three-minute thesis

NAME: Lucie Klus
RESEARCH CENTER: Tampere University (Finland) and University Jaume I (Spain)
SUPERVISORS: Jari Nurmi, Elena Simon Lohan and Carlos Granell
DATE DEFENDED: 27/04/2023
THESIS TITLE: From compression of wearable-based data to effortless indoor positioning

In 2022, almost 500 million wearables were sold and shipped worldwide, and the exponential growth in the market in smart watches, fitness bands, e-textiles and smart glasses is expected to continue. Wearable devices have become ubiquitous, enhancing daily life through intelligent features and services. However, their effectiveness is constrained by their compact size and limited battery capacity, which mean that energy efficiency is a key requirement.

In my doctoral dissertation, I concentrated on enhancing the energy efficiency of IoT and wearable devices by incorporating lossy compression techniques into sensor-based time-series data and indoor localization paradigms. Connected devices handle vast amounts of diverse data, and optimizing algorithms leads to valuable computational and storage savings. Rather than utilizing powerful machines and complex algorithms, therefore, my research had to take on the challenges of maximizing the effectiveness of the device while minimizing use of algorithmic resources.

For time-series data, the degree of information loss caused by compression is as important as the energy or time that the algorithm needs to execute. On the other hand, information loss is not a relevant metric in the scope of indoor positioning, where the ability of the data to ensure seamless and accurate localization is the critical objective. However, indoor positioning does require localization methods that don't rely on Global Navigation Satellite System (GNSS) modules, as the latter can only ensure high-accuracy positioning when several satellites are directly observable.

Data compression that doesn't drain the battery

During my doctoral research, I developed novel compression techniques, a new performance metric, and the means to objectively evaluate a single algorithm in a multitude of environments, incorporating aspects of numerous scientific fields. I also addressed several of the challenges thrown up by using GNSS-free localization methods. Overall, my research aims to enable the acquisition of high-quality data from wearable and IoT devices while significantly prolonging battery life.

In addition to using a vast number of publicly available datasets to ensure the validity of the methods proposed, I performed my own site survey in which I created a multi-device dataset while on a six-month research visit to Jaume I University in Spain. With this hands-on experience, I discovered new research opportunities, developed my personal capabilities, and identified challenges which hadn't previously been addressed. Although my doctoral journey has now come to an end, I am certainly not done with research and all the excitement it offers.

Lucie's PhD supervisor, **Jari Nurmi**, commented: 'Lucie's dissertation proposed many innovative algorithms and techniques for compressing data on resource-constrained wearable devices and validated these through experimentation. Compression is particularly important for reducing the computational load and communication needs on devices with limited resources, thus saving energy and extending their battery life. The data included time series sensor-based data, as well as location-dependent data. The 25 public datasets she used in evaluating the algorithms were extensive and diverse. Furthermore, she collected a dataset of her own that is now published as open source, thus contributing to reproductivity in the research community.'

HiPEAC conference

17-19 January 2024

Munich

computer architecture · machine learning · compilation · cybersecurity · energy efficiency · open source · edge · cloud · internet of things · high-performance computing · cyber-physical systems

hipeac.net/conference

#HiPEAC24

Funded by the European Union