

INSIDE


magazine



July 2025 — Issue 10

In this issue:

- **Computing with the cosmos**
- **AI in automotive**
- **Driving impact beyond research**




INSIDE
Industry Association

In this issue

July 2025 — Issue 10

- 04 — Computing with the cosmos
- 10 — Past, present and future of RISC-V and its role in Chips-IT
- 15 — INSIDE Connect 2025
- 16 — AI in automotive
- 22 — The micro-nano facility at the Fondazione Bruno Kessler
- 26 — Pioneering the future of mobility
- 28 — Driving impact beyond research
- 32 — Crossing Europe's valley of death
- 34 — Cubit Innovation Labs
- 37 — ENSURE Workshop
- 38 — AIMS5.0
- 40 — CEI-Sphere



Computing with the cosmos



AI in automotive



Driving impact beyond research

Dear reader,

In this issue, we explore the frontiers of computation, the vibrancy of our community, and the evolving landscape of science, engineering, and collaboration.

We open with a journey to the very edge of physics, which investigates the radical idea of harnessing gravitational waves for computation. From Einstein's curved spacetime to speculative gravitational logic gates, the article examines how recent theories, including Tselentis and Baumeler's work, might one day turn ripples in spacetime into information processors.

We then move to the microarchitecture of processors, with a timely reflection on how open standards are reshaping Europe's strategic position in chip design. The story traces the evolution of a promising RISC-V technology stack and shows how new initiatives are aligning efforts from research to production.

In AI, we focus on how it is transforming automotive development through deep integration of engineering automation. AI-driven tools now generate real-time, safety-critical code and automate testing, accelerating delivery and improving consistency. This enables modular, scalable architectures tailored to the demands of autonomous and connected vehicles.

Zooming into the device level, another article showcases FBK fabrication facility that has evolved into a complete pilot line. By offering a cleanroom platform for advanced sensors, microstructures, and silicon integration, it stands as a model for bridging early-stage research and industry-grade manufacturing.

Manufacturing itself is undergoing a shift toward human-centric and sustainable systems. AIMS 5.0, a large-scale initiative highlighted in this issue, focuses on embedding AI within production processes, not only to improve efficiency, but to advance resilience, transparency, and energy responsibility across supply chains.

Also in the projects sector, CEI-Sphere is building Europe's digital resilience by aligning Cloud-Edge-IoT technologies with real-world market needs. Through sector-specific Market Briefs and the CEI Market Talks series, it connects technological innovation with strategic insights for mobility, energy, and manufacturing. As it supports the continent's

Large-Scale Pilots, CEI-Sphere is becoming a key voice in Europe's digital transformation

The New Member Focus section introduces fresh voices to the INSIDE community.

From software foundations to mobility platforms, we visit Apex.AI, a company pioneering middleware systems for deterministic, safe, and scalable automotive computing, offering a glimpse into how embedded intelligence is being re-engineered for the real world.

We then step into the strategic arena of research valorization with Innovation Dis.Co. A team dedicated to dissemination and exploitation reveals how real-time analytics, clear narratives, and actionable strategies can turn project results into long-term societal and industrial impact.

With Quintauris we explore the so-called "valley of death" in Europe's semiconductor landscape, outlining systemic issues that block innovation from becoming product. With sharp analysis and actionable proposals, it raises the question: how do we shift from invention to adoption, before momentum is lost?

Finally, Cubit presents a regional innovation lab bringing us back to applied research. From wireless systems to air quality sensing, this lab exemplifies agile collaboration between academia and industry, proving that innovation is just as much about people and partnerships as it is about technology.

What ties these stories together is a shared orientation toward transformation, of knowledge into systems, of research into application, of vision into structure. Whether exploring how we compute, manufacture, move, or collaborate, each contribution reflects a deliberate shift toward shaping - not just anticipating - the future. The future, as always, begins with asking the right questions.

Paolo Azzoni
Secretary General



Technology Frontiers

Computing with the cosmos

Can gravitational
waves power the
computers of
tomorrow?



Paolo Azzoni

Recently, some scientists have proposed the exotic idea that the existence of gravitational attraction is proof that our universe is nothing but a simulation (remember *The Matrix* movie?)^{1, 2, 3}. This theory suggests that gravity itself is a computational artifact, implying that it takes a computer to make gravity. Luckily, this notion was recently challenged by Italian astrophysicist Franco Vazza⁴, who demonstrated that simulating even a small portion of the universe would require more energy than is physically available, making the simulation hypothesis practically impossible.

But what if we flipped the question: could gravitational waves themselves be harnessed for computation or information processing? In other words, can we envisage a gravitational wave computer in the future? This concept sits at the far edge of unconventional computing paradigms (beyond quantum or optical computing) blending general relativity with information science.

In this article, I want to be more exotic than usual (we are in summer ...) and explore with you the scientific foundations of gravitational-wave-based computing, current research and speculative models, the challenges, and how this idea compares to other computing approaches.

Space-time: the fabric of physical reality


Einstein's theory of general relativity explains that space-time is curved by mass. This curvature is the new definition of gravity. The trampoline analogy explains with an example how massive objects curve space-time and how small objects move along this curvature (see the figure). Gravitational lensing shows that even the path of light can be bent in curved space-time - just like an optical lens. But Einstein's theories of special and general relativity are contradictory: is space-time an unchanging expanse, or can it be warped in ways that affect a signal travelling through it? According to special relativity it is static, while the theory of general relativity reveals something completely different: in this framework, massive objects make space-time dimple and curve, like when a ball is dropped onto a taut sheet, which could change the path of a signal moving nearby.

Albert Einstein's general theory of relativity revolutionized our understanding of gravity by describing it not as a force but as a curvature of spacetime caused by mass and energy⁵.

Accelerating massive objects generate gravitational waves, which can be conceived as ripples in spacetime that propagate at the speed of light. These waves were first predicted by Einstein in 1916 and remained undetected until September 14, 2015, when the Laser Interferometer Gravitational-Wave Observatory⁶ (LIGO) made the first direct observation of gravitational waves from a binary black hole merger. This landmark discovery confirmed a major prediction of general relativity and opened a new window into the cosmos.

Gravitational waves carry information about their cataclysmic origins and travel virtually unimpeded through matter, curving even the path of light and providing insights into phenomena that are otherwise invisible.

In the framework of relativity, the causal structure of spacetime dictates the cause-effect relationships between events, constrained by the speed of light. This structure is represented by light cones, which



define the possible influence of events in spacetime⁷.

While special relativity treats spacetime as a static, flat expanse, general relativity reveals its dynamic nature, where mass and energy can warp spacetime, altering the trajectories of objects and signals. This interplay between the geometry of spacetime and the distribution of mass and energy is central to our understanding of the universe's structure and behaviour.

Foundations of gravitational-wave Computing

At its core, a gravitational wave (GW) is a traveling distortion in spacetime. Unlike electrons in circuits or photons in fibres, gravitational waves are *pure spacetime vibrations* that, once generated, propagate

In a computing or communication context, a signal encoded in gravitational waves could be received anywhere in the universe with negligible loss, something impossible for electromagnetic signals which are blocked or slowed by materials (even light is slowed by media like fiber to ~70% of c).


at constant light speed and are hardly absorbed or scattered by matter. This unique property means GWs could, in theory, act as **ultra-reliable information carriers**, able to traverse stars, planets, or any medium without attenuation. Indeed, astronomers note that GWs from distant cosmic events have travelled billions of years essentially unaltered.

Another foundational aspect is that GWs, being solutions of Einstein's field equations, interact with spacetime geometry and potentially with quantum fields. This has led to theoretical considerations of **information in spacetime**. For example, gravitational waves carry energy and information about the motion of their source masses. In principle, any dynamic gravitational field can be modulated, by changing the motion of masses, to encode data in the wave's amplitude, frequency, or phase. A simple thought experiment: a spinning dumbbell emits weak gravitational waves; if you could vary its rotation speed, you would frequency-modulate those waves and transmit a message. Thus, the physics *allows* GWs to be used as an information-bearing signal. In fact, gravitational waveforms detected by LIGO already **encode information** about black hole masses and spins.

The speculative leap is to go from passively reading cosmic signals to actively generating and processing information with GWs.

However, using GWs for computing means treating spacetime distortions as computational states. This ventures into the territory of **relativistic computation** and even quantum gravity. Some theoretical computer scientists have studied exotic general-relativistic models where gravity or spacetime structure enables new computation modes (e.g. "Malament–Hogarth" spacetimes that allow, in theory, non-Turing computations by exploiting relativistic time dilation⁸).

A GW-based computer would essentially use general relativity itself as the computing substrate. The concept also overlaps with quantum gravity ideas: for instance, physicist



Massive bodies like the Sun and Earth warp this fabric, influencing the motion of other objects and of light itself. The trampoline analogy provides us an example easy to understand.

Lucien Hardy coined the term “*quantum gravity computer*”, referring to a device where quantum mechanics and general relativity jointly inform the computation⁹. In such a machine, information might be encoded in quantum states embedded in spacetime, and logic gates would involve moving enough mass-energy to measurably curve spacetime or generate gravitational waves. The theoretical elegance is that combining quantum superposition with spacetime warping might perform feats impossible for standard Turing machines – perhaps even solving non-computable problems by utilizing relativistic effects.

This is highly speculative, but it provides a conceptual foundation: if we had fine control of gravity at quantum scales, we might leverage it for unprecedented computing power.

occur by sending those waves through spacetime and having them interact. The appeal of this model is that communication of results could be instantaneous across cosmic distances (at light speed with no delays due to matter). For example, a gravitational logic gate could conceivably send its output as a gravitational ripple to any location without needing physical wires or optical links. Furthermore, Tanaka speculates that a gravitational computer might access past and future states of the universe by manipulating spacetime, effectively blurring the line between computation and time travel. While this ventures into science fiction, it stems from the relativistic fact that gravitational disturbances can affect remote points in spacetime.

A gravitational computer might “search across past, present, and future states of the universe” by cleverly encoding queries into gravitational disturbances.

Existing speculative theories

Current theories are largely speculative or at best at the proposal stage. One prominent vision was outlined by futurist Shoichi Tanaka, who imagines a distant-future scenario where **spacetime itself is the computing medium**, a “meta-timespace computing” paradigm. In this vision, information would be encoded into gravitational waves, and computation would

This dramatic idea underscores how gravitational-wave computing, if ever realized, could fundamentally alter notions of causality and information flow.

Another theoretical proposal comes from efforts to merge quantum computing with gravity. The **relativistic quantum**

gravity computer mentioned above is one such framework: it envisions qubits tied to spacetime such that operations involve oscillating matter and gravitational wave emission¹⁰. One patent application describes logic gates that require moving a certain mass threshold during operation so that each gate’s action perturbs the local spacetime metric with gravitational waves. If the perturbation is significant enough, it triggers a collapse of the quantum state (an idea inspired by Penrose’s hypothesis that gravity can induce wavefunction collapse¹¹). In essence, these gates would perform logic by balancing quantum effects against gravitational effects. The proponents argue such a “*quantum gravity computer*” could solve problems beyond the reach of standard quantum computers because it isn’t bound to a fixed causal structure. Though still hypothetical, this represents a concrete (if extremely ambitious) model: a computing device where **each logical operation emits and detects gravitational waves in a controlled fashion**, using them to influence quantum bits. It’s a marriage of qubits and curved spacetime, expected to be more powerful than Turing machines.

Beyond these grand visions, scientists have considered more classical uses of gravitational waves in information processing. For example, proposals for **gravitational wave communication networks** are actively being explored at the theoretical level. A recent comprehensive review highlights how one might modulate gravitational waves to carry data, discussing techniques like amplitude modulation via astrophysical oscillators or frequency modulation using phenomena like ultralight dark matter clouds¹². One scheme suggests using a pair of superconducting resonators: by oscillating the mass density in one superconducting block, a high-frequency GW signal could be generated and picked up by a second superconducting detector, forming a basic transmitter-receiver pair. In theory, this could be a **gravitational-wave analogue of radio communication**, just with vastly different hardware: one needs a generator (mass in motion), a modulator (to encode bits onto the wave, perhaps by varying motion or mass distribution), and a sensitive detector to read the faint ripples ... not easy to develop.

Finally, there is an analogy with **neutrino-based computing/communication** proposals¹³, another exotic paradigm. Like gravitational waves, neutrinos hardly interact with matter and can travel through

planets. Researchers have demonstrated rudimentary neutrino communication through rock and even speculated on neutrino-based quantum computers using neutrino as qubits¹⁴. Gravitational wave computing is even more extreme: instead of particles, it uses spacetime itself as the carrier. But the motivations overlap: in both cases, the goal is to exploit a medium that can traverse any obstacle and perhaps access regimes standard signals cannot.

A concrete leap forward?

In addition to their unique propagation properties and information-carrying potential, GWs open a profound avenue for rethinking how computation may emerge from spacetime dynamics itself (special relativity, static – general relativity, dynamic). Recent work by Tselentis and Baumeler¹⁵ contributes significantly to this new perspective by proposing a mathematical framework to assess whether spacetime in a given region can be considered fixed or subject to distortion. Their work bridges a critical gap between general relativity and information theory, showing that the ability to reorder messages between observers can serve as a test of spacetime warping.

This has direct implications for GWs-based computing: if manipulating spacetime allows one to alter causal structures, such that signals can be received in a different temporal order depending on the distortion, then computational processes could, in principle, be encoded in the causal topology itself. Tselentis and Baumeler's framework formalizes these concepts using communication protocols between 3 people (like Alice, Bob, and Charlie), evaluating whether spacetime deformation has occurred simply by observing message arrival patterns.

This formalism aligns with the speculative leap in gravitational computing, shifting from passively observing gravitational waves to actively engineering spacetime geometry to process information. It reinforces the possibility that future gravitational computing models won't just rely on high-energy mass movements to generate waves, but might encode logic through dynamic causal configurations, suggesting spacetime curvature is not just a background for computation: it is the substrate.

While many current models of GWs computing are theoretical and speculative, the contribution by Tselentis and Baumeler adds new weight

to the notion that spacetime dynamics can be systematically linked to information flow. Their work departs from previous models that often rely on extreme or hypothetical constructs, like quantum gravity gates or spacetime machines, to propose a testable, observer-independent method for detecting spacetime manipulation via communication patterns between distributed agents.

This insight opens up a radically different type of theoretical model: instead of engineering specific gravitational wave sources to encode information, one might analyze and program computation through the manipulation of causal paths in a spacetime manifold. In their setup, message delays and reorderings caused by gravitational distortions become logical signals, enabling a form of "protocol-based spacetime computation"

This approach could lead to entirely new proposals: for instance, causal-layered computing, where algorithms are defined not in Boolean logic but in relativistic message-order permutations. Such a paradigm redefines what it means to "compute," and places gravity at the heart of logic itself.

Such a model doesn't require speculative components like exotic particles or unproven materials. Instead, it draws directly from communication theory, general relativity, and the causal structure of spacetime. Importantly, it frames computation not as a sequence of operations inside a machine, but as a distributed process governed by the geometry of the universe.

Major challenges

Computing with GWs presents a long list of critical challenges, including:

- 1 Extremely weak interaction: GWs barely interact with matter, producing distortions as small as 10^{-18} meters, making detection

and use for computing extraordinarily difficult without massive, ultra-sensitive equipment.

- 2 Generation and energy requirements: creating meaningful GW signals requires astronomical energies or massive moving objects. Lab-scale generators produce waves too weak to detect, and current technology lacks efficient methods to amplify them (spinning ultradense rotors, high-frequency vibration of large masses, resonant cavities, and even beam collision devices)¹⁶.
- 3 Control and modulation^{9, 17}: GWs cannot be focused, shielded, or switched like light or electricity. Their broad propagation and limited polarization make it hard to isolate and control them for logic operations or signal routing.
- 4 Decoherence and precision: if quantum effects are involved, gravitational computing inherits challenges like decoherence and noise sensitivity. Achieving necessary environmental isolation for stable operations may require deep-space conditions.
- 5 Theoretical unknowns: without a complete theory of quantum gravity, key mechanisms remain speculative. Issues like time loops, frame dragging, and causal paradoxes complicate the idea of using spacetime as a reliable computing medium.

Potential future applications

Given these challenges, this technology is likely many decades away (if not impossible with known physics). However, some potential applications can be envisaged, including:

- Planet-scale communication: GWs could transmit data through planets or stars with no attenuation or interference, enabling ultra-secure, jam-proof communication, even in environments where radio fails.
- Deep space networks: GWs could form the backbone of a solar system (if not galactic ...) internet, offering direct, line-free signaling across cosmic distances, ideal for solar system and interstellar missions or off-world colonies.
- Computing in extreme environments: unlike electronics, GW-based systems might function in high-radiation, high-magnetic, or high-temperature zones (e.g., near nuclear reactors) without signal degradation.
- Ultra-precise timekeeping and navigation: pulsed gravitational waves could serve as universal, dispersion-free time signals,



enabling synchronized processors and gravity-based positioning across vast distances.

- Astrophysical computers: natural GW sources like pulsars might act as components in large-scale analog computers, where their interactions encode computations at cosmic scale, more physical than digital.
- Trans-temporal signaling: in principle, GWs could exploit spacetime geometry to transmit information across time, hinting at future devices capable of time-shifted logic or feedback from the future.

Many of these applications read like science fiction scenarios precisely because gravitational computing blurs the line between computation and fundamental physics but, even if we never achieve these applications, thinking about them can inspire new questions in physics and inform other areas.

Gravitational computing vs. the frontier: how it measures up

Gravitational wave computing could compete with other unconventional paradigms that challenge traditional silicon/optical-based architectures. How does it compare?

- Quantum computing: both aim to surpass classical limits. Quantum computing exploits superposition and entanglement, while gravitational computing could leverage relativistic effects like time dilation. But GWs systems face extreme decoherence, lack scalability, and are purely theoretical, though future hybrid models (e.g. gravity-assisted qubits) are conceivable.
- Optical/photonics computing: photonic systems are fast, scalable, and well-developed, using light in compact chips. GWs travel without scattering but need massive setups. Unlike optical systems compatible with silicon tech, gravitational computing lacks a physical substrate, because its “platform” is spacetime itself.

- Biological/neuromorphic computing: biological systems are slow but massively parallel and self-healing. Gravitational computing is the opposite: precise, fast, but rigid and sensitive to noise. Programming would require orchestrating spacetime itself, with no current frameworks or languages.
- Other wave-based systems: wave-based models (e.g. spin waves, solitons) compute via interference and collisions. GWs computing could do the same, interfering spacetime ripples instead of material waves, but suffers from weak signal strength and impractical scale, limiting short-term feasibility.

Computing at the edge of spacetime

Gravitational wave computing may be the most mind-bending idea in the computing frontier, a paradigm with cosmic reach but firmly out of technological capabilities, at least for now. Where quantum excels in speed, optical in bandwidth, and biological in efficiency and parallelism, gravitational computing teases a niche of its own: unhindered transmission through matter and perhaps, someday, manipulation of space and time itself.

But this isn't just about building machines, it's about reimagining what *computation* even means. Gravitational waves don't travel along wires; they ripple through the very structure of reality. To harness them, we may need to rewrite the foundations of computer science, embracing models that are analog, asynchronous, and curved like spacetime itself.

As with quantum computing decades ago, even thinking seriously about gravitational computation pushes the boundaries of physics, logic, and engineering. And perhaps, in that act of imagining, we'll discover not just a new kind of computer

but a deeper understanding of the universe it would compute within.

And perhaps the first true leap has already happened: the recent work by Tselentis and Baumeler offers a concrete framework connecting information flow and spacetime variability, a rare bridge between abstract theory and the physical structure of causality. If their insights prove fertile, gravitational computing may no longer be just a thought experiment, but the start of a radically new information paradigm.

¹ Bostrom, N. (2003). Are You Living in a Computer Simulation? *Philosophical Quarterly*, 53(211), 243–255.

² <https://arxiv.org/abs/2404.16050>

³ <https://doi.org/10.1063/5.0173275>

⁴ <https://doi.org/10.3389/fphy.2025.1561873>

⁵ <https://ijrar.org/papers/IJRAR19D5434.pdf>

⁶ <https://arxiv.org/abs/1602.03837>

⁷ <https://modern-physics.org/causality-in-relativity/>

⁸ <https://arxiv.org/pdf/1105.0047>

⁹ Relativistic quantum computer / quantum gravity computer, US Patent US20220366289A1

¹⁰ Relativistic quantum computer / quantum gravity computer, US Patent US20220366289A1

¹¹ DOI: 10.1039/D4CP02364A

¹² <https://arxiv.org/html/2501.03251v1>

¹³ <https://physicsworld.com/a/neutrino-based-communication-is-a-first/>

¹⁴ DOI: 10.13140/RG.2.2.28062.14409

¹⁵ <https://journals.aps.org/prd/abstract/10.1103/PhysRevD.111.052211>

¹⁶ <https://arxiv.org/html/2501.03251v1>

¹⁷ <https://arxiv.org/abs/2109.01535>

Past, present and future of RISC-V and its role in Chips-IT

RISC-V



Davide Rossi



Yvan Tortorella



Cristina De Luca



Luca Benini

We are living in a transformative computing era. The surging demand for automation and ubiquitous connectivity, along with the explosive growth of artificial intelligence (AI) have radically expanded the boundaries of what is technologically possible. This era has fuelled an immense demand in the semiconductor industry, initially focused on large datacenter chips for AI acceleration, but now also driven by the proliferation of IoT sensors, wireless connectivity, and embedded “physical AI” systems crucial for robots, autonomous vehicles and industrial automation. Further accelerating this demand are the increasing needs for mobile bandwidth and vehicle electrification.

A brief history of reduced instruction set computing

The RISC-V Instruction Set Architecture (ISA) originated in 2010 at the University of California, Berkeley, as a purely academic initiative. Since then, RISC-V has continuously gained traction, not only in academia but also across industry and public institutions, revolutionising the design of digital architectures. Much like what Linux did for operating systems-transforming a university project into the foundation of a vast and collaborative open-source ecosystem, RISC-V is enabling a new paradigm for processor design that encourages broad participation and customisation.

Notably, the design flexibility offered by RISC-V architecture has catalysed a revolution across the semiconductor sector and numerous related industries, including automotive, space and aviation, security, IoT, wearable devices, robotics and high-performance computing.

RISC-V is an open standard instruction set architecture based on the well-established reduced instruction set computing (RISC) principles. Unlike proprietary ISAs such as x86 or ARM, RISC-V is provided under open licences, free from royalty fees or restrictive legal conditions. It is built around a minimal base instruction set, to which optional standard or custom extensions can be added to tailor the processor's capabilities to specific application needs. What further distinguishes RISC-V is the way it evolves: the specification of new

instructions and extensions is managed by RISC-V International, a non-profit organisation headquartered in Zurich (for international neutrality). Its governance model is inclusive and collaborative, enabling members-from startups and universities to large technology companies from most countries around the world to participate in the definition and ratification of the ISA through a transparent, consensus-driven process. This democratic approach ensures that RISC-V reflects the needs of a diverse global community, while avoiding the constraints and vendor lock-in typical of closed architectures.

In parallel to its Californian counterpart at Berkeley, which focused primarily on specification of the instruction set architecture (ISA), the PULP (Parallel Ultra Low Power) project has led the development of RISC-V Intellectual Property (IP) components in Europe. The PULP project was born as a collaborative effort between the University of Bologna and ETH Zurich. Initially based on the OpenRISC architecture (a former open-source ISA developed in the early 2000s), PULP aimed to explore energy-efficient multicore processors for deeply embedded applications. After attending the first RISC-V workshop in Monterey in 2014, the PULP project members decided to transition to the RISC-V ISA, pushed by the recognition that RISC-V offered a rapidly growing global community and ecosystem that other open-source projects like OpenRISC lacked.

Since then, PULP has embodied the spirit of RISC-V by delivering a wide set of open-

source IPs, software stacks and several full system-on-chip (SoC) architectures tailored for IoT, embedded computing, and later high-performance computing and AI. It has produced dozens of open SoCs, IP blocks and verification platforms, making it one of the most influential RISC-V research platforms in the world, with a large number of users not only in academia but also across industry, thanks to its liberal, permissive licensing model. The success of PULP has positioned Europe as global leaders in the research related to the design of customisable, high-efficiency RISC-V cores, and has laid the groundwork for national and continental ambitions in open digital hardware innovation.

This leadership is reflected in the PULP team's strategic involvement, particularly from the University of Bologna and ETH Zurich, in key European initiatives aimed at strengthening technological sovereignty through open and customisable compute platforms. Examples include EPI, TEP and DARE, which aim to enhance European autonomy in high-performance computing by developing processors and accelerators, along with the complete software stack needed for RISC-V-based HPC applications. Projects like TRISTAN and ISOLDE focus on embedded and industrial domains, including IoT and real-time control, while RIGOLETTO and HAL4SDV target the automotive sector, respectively addressing next-generation hardware platforms and the supporting software stack for safe and secure Software Defined Vehicles.

The ambition in the Italian industrial ecosystem

In this context, the Italian government has launched Chips-IT, the national centre for semiconductor research and innovation, as part of the broader European Chips Act strategy. Its mission is to strengthen Italy's capabilities in microelectronics by creating a distributed infrastructure that supports advanced research, pre-industrial prototyping, and technology transfer to industry. Within Chips-IT, Bologna has been selected as the home of the "Digital Design and Open Hardware" line, precisely because of the internationally recognised leadership of the University of Bologna in open, energy-efficient processor design, clearly demonstrated through its contributions to the PULP platform and RISC-V ecosystem.

The ambition of this activity within Chips-IT is to build upon the solid foundation of research and the industrial ecosystem developed

through initiatives like PULP, by leveraging the more advanced infrastructure and dedicated resources available in a national research center, typically more robust and industry-oriented than those of individual universities. The goal is to use the PULP open-source IPs as a catalyst to foster a broader Italian ecosystem around digital design and open hardware, actively involving other universities, research institutions and industrial players. This initiative also aims to grow a national community of expertise and collaboration, leveraging open-source best practices to help address a critical gap: the shortage of electronics and chip design engineers in Italy, which increasingly limits the competitiveness and innovation capacity of the country's industrial landscape.


Chips-IT aims to further develop and mature the technologies created by open source projects such as the PULP project, making them more accessible and ready to be adopted by Italian companies. This includes working alongside industrial partners to co-design and validate new solutions across several key strategic sectors such as IoT, wearable devices, robotics, aerospace, automotive and high-performance computing. Through these efforts, Chips-IT not only contributes to reinforcing Italy's technological sovereignty, but also builds a solid bridge between open academic research and industrial innovation-ensuring that the success story of RISC-V continues to grow within a strong, collaborative and future-oriented national and international framework as well as a strong industrial ecosystem.

Industrial adoption and momentum in the RISC-V ecosystem

The growing industrial interest in RISC-V is already visible in the vibrant European ecosystem, where several companies, ranging from research-driven startups to multinational joint ventures, are embracing the open ISA to build differentiated, domain-specific computing platforms.

Among the pioneers of the RISC-V movement is SiFive, the California-based company founded by the original creators of the ISA. SiFive has played a foundational role in making RISC-V commercially viable, offering a scalable portfolio of customisable cores now used in applications ranging from microcontrollers to high-performance computing. Its technology forms the basis for many academic and industrial developments worldwide, including in Europe, where its cores are often integrated into larger system





designs or used as reference platforms for innovation.

On the European front, the fabless startup Axelera AI exemplifies how RISC-V enables architectural freedom in the context of edge AI. The company has developed a novel acceleration platform that combines in-memory computing with RISC-V-based control logic, creating a tightly integrated, low-power system for real-time inference in embedded devices. RISC-V's openness and modularity have been instrumental for Axelera in fine-tuning hardware behaviour to meet stringent latency and energy constraints—something difficult to achieve with more rigid proprietary cores. Axelera's trajectory also highlights the growing investor confidence in RISC-V-driven innovation: to date, the company has raised over \$200 million in funding, positioning itself as one of the best-capitalised European deep-tech startups in the AI hardware space.

Another landmark initiative is Quintauris, a newly formed European joint venture aiming to develop RISC-V-based processors for the automotive, industrial and broader embedded markets. Initially established by Robert Bosch GmbH, Infineon Technologies AG, Nordic Semiconductor, NXP Semiconductors and Qualcomm Technologies, the venture was recently joined by STMicroelectronics, further solidifying its industrial weight. By uniting some of the most influential semiconductor players under a common RISC-V strategy, Quintauris embodies Europe's ambition to reduce dependence on foreign IP monopolies and assert control over critical compute infrastructure. The consortium aims to accelerate the availability of robust, production-grade RISC-V solutions that meet the demanding requirements of safety, security and long-term support, particularly in automotive and industrial domains.

It is important to stress that hardware-focused RISC-V initiatives are also supported with large critical-mass international software-centric industry alliances. As an example, the RISC-V Software Ecosystem (RISE) flagship project is a collaborative effort led by industry leaders with a mission to accelerate the development of open source software for the RISC-V architecture. RISE is supported by more than 20 companies from across the world (USA, ASIA, EU), including industry leaders like Qualcomm, Samsung, Mediatek, Intel, AMD, NVIDIA.

Together, these examples reflect the depth and maturity of the RISC-V ecosystem:

from global IP leaders like SiFive to agile innovators like Axelera AI and large Europe-centric industrial alliances like Quintauris. RISC-V is enabling a new paradigm of innovation in the design of digital ICs and SoCs, one rooted in openness, customisation, and sovereignty—values that are increasingly central to Europe's and Italy's technological strategies.

Market trend and expected impact

The RISC-V market is evolving with remarkable speed, driven by new entrants who are significantly advancing innovation, especially in powerful CPU IP types. This dynamic landscape couples with multi-dimensional challenges such as semiconductor shrinking, the need for enhanced energy efficiency and increasing computational complexity, profoundly impacting AI silicon development and multi-core System-on-Chip (SoC) design. These factors, in turn, affect design and IP licensing revenues.

Despite this complexity, the HSD group's forecast predicts substantial growth for the RISC-V market, fuelled by escalating chip demand. The market is projected to expand significantly from \$6.1 billion in 2023, growing by 276.8% by 2030. Based on a RISC-V market trend and prediction from 2024¹, market share distribution is set to grow: industrial applications are predicted to remain stable (from 2.6% in 2023 to 2.7% in 2030), while networking is projected to experience significant growth (from 7.1% in 2023 to 18% in 2030). Automotive's share is also expected to climb (7.4% to 9.9%), whereas the consumer segment's share could maintain relative stability (33.8% to 34%). This anticipated expansion is largely seen as being driven by the rising demand for customized chips. A key driver for this expansion is the rising demand for customised chips.

Furthermore, the forecast indicates a crucial shift in revenue models: RISC-V market revenues based on CPU IPs are expected to surpass those from licences by around 2027, with royalties projected to reach 1.5 times licence revenues by 2030. This suggests license revenues are approaching a saturation point due to vendor competition and bulk sales. RISC-V's inherently open nature is a distinct advantage, as it fosters innovation from multiple companies, each focusing on diverse markets and applications. This collective approach accelerates and diversifies advancements beyond what any single company could achieve¹.


The evolving computing landscape, driven by automation and AI, necessitates significant advancements across the semiconductor value chain. Open-source ISAs like RISC-V require further maturation of their software stacks to fully leverage hardware advantages for Domain-Specific Architectures (DSAs). Simultaneously, the paradigm shift from Systems-on-chip to System-in-Package (SiP) and chiplet (2.5D) integration demands open standards not only for on-chip but also for chiplet-to-chiplet interfaces.

Electronic Design Automation (EDA) is also undergoing a deep transformation: tools need to support the design of domain specific architecture, expanding their scope from SoC to SiP solutions, managing complex, multiphysics system-level design and verification challenges across multiple chiplets. A new form of co-design is emerging, where design automation will be shaped by domain-specific AI-driven enhancements. In this context, open-source EDA can play a key role, as it has the potential to break decades-old boundaries in design optimisation: open-source design automation tools can be tailored with domain-specific optimisers (e.g. exploiting the architectural regularity of systolic datapaths) pushing optimisation more aggressively.

Ultimately, McKinsey Direct² suggests that Domain-Specific OEMs (e.g., IoT, automotive) must shift from general-purpose compute approaches to deeply understand their specific workload needs, exploring optimised chip architectures, and potentially engaging directly with DSA startups beyond traditional supply chains.

Takeaway

The true power of the RISC-V ecosystem transcends mere cost savings. While its potentially lower licensing fees are often cited as an advantage over traditional Instruction Set Architectures (ISAs) like ARM, RISC-V's profound value lies in the broader, more transformative benefits it delivers directly to the global design community based data models.



¹ Shivangi Kharoo, "RISC-V Market Trends And Predictions Till 2030", June 2024

² "Domain-specific architectures and the future of compute", July 18, 2023 | McKinsey Direct

INSIDE Connect 2025

3-4 September
Malaga, Spain

The INSIDE Connect 2025 event — from 3 to 4 September — will give you a unique opportunity to re-engage, collaborate, and shape the future of our community.



For more information and programme got to <https://inside-association.eu/inside-connect-2025/> or scan the QR code.



Register here or scan the QR code.

It's time to connect!

We warmly invite our Members and personally invited guests to INSIDE Connect 2025 event in Malaga, Spain! We are undergoing a profound transformation in which the geopolitical landscape, technological innovation, sustainability imperatives, technological and economic competitiveness, and evolving industrial priorities are reshaping the future at an unprecedented pace.

INSIDE Connect 2025 is where our interdisciplinary, inclusive, and forward-thinking community comes together to anticipate what lies ahead, defining a collective vision that strengthens our impact, influence, and identity in a rapidly changing world. This is where connections become direction. And where tomorrow begins — together.

What to expect?

Positioning our community for the next European MFF

As the next EU Competitiveness Fund and Framework Programme take shape, our community must come together to define its strategic position and priorities in Europe's digital future.

The future of digital innovation

In times of fast technological changes, this event will align key stakeholders around shared priorities and drive tangible next steps for 2025–2026, turning ambition into action.

A space to connect

Beyond strategic alignment, this is a place to connect — building new relationships, forging partnerships, and empowering the trust needed to steer the next wave of innovation across Europe's digital ecosystem.

Workshops & road mapping

Interactive workshops and road-mapping sessions will drive deep dives into key thematic areas, enabling participants to co-define strategic directions and shape actionable priorities. These sessions are where ideas crystallize into roadmaps, guiding our collective efforts and investments through 2025 and 2026.

Project Showcases & Exhibition

See innovation in action. Selected projects will present results and exhibit tangible outcomes, offering visibility, fostering knowledge exchange, showcasing impact, and sparking new collaboration across the digital innovation ecosystem.

AI in automotive

The engineering challenge



Petri Solanti

New automotive products apply several complex artificial intelligence algorithms related to autonomous driving and assistant systems. Because of the high algorithm complexity and limited computing resources, the implementation must be carefully optimised. The complexity of these systems exceeds the capabilities of traditional design methodologies. The growing interdependence of software and hardware requires a new design methodology with a strong focus on abstract system analysis, multi-disciplinarity and multi-context capabilities.

A new engineering methodology was developed to solve the design challenges of complex cyber-physical systems. It allows the system architects to explore different implementation options with reasonable accuracy on an abstract level and refine the system architecture, HW/SW partitioning and HW implementation to the level that can be handed over to the domain engineers. Combined with High-Level Synthesis, this methodology provides a seamless flow from top-level system requirements down to RTL-implementation of tailored HW accelerators on FPGA or System-on-Chip.

The challenge of the edge

By the definition of edge device, vehicles are among the most complex. A car is an autonomous electronics platform without a connection to a power supply or wired communication network. Despite that, more and more Artificial Intelligence (AI) functionalities are being added to modern vehicles to enable assisted or fully autonomous driving.

AI algorithms are usually developed and tested in data centres that have practically unlimited computing, memory and data storage resources, high-speed wired data communication and a strong uninterruptible power supply to enable continuous high-performance computations.

In a vehicle these capacities are limited, which constrains the selection of possible AI algorithms and how they can be executed. Additionally, the requirements like response time, enclosure size, power consumption, network bandwidth or thermal dissipation constraints limit the design options and may even prevent use of the selected algorithm.

Automotive platforms are standardised at company level, and the manufacturers use the same Electronic Control Units (ECU) in the different car models as much as possible to save costs. This leads to another problem for AI engineers, who must find algorithms that can be executed on the given platform. The lead time to develop a new ECU can be several years, if a new dedicated System-on-Chip (SoC) is needed. Typically, a new ECU development project requires multiple subcontractors to work together, which causes additional communication and data protection problems.

Developing optimal implementations of AI systems in an automotive environment requires a new engineering methodology that enables design teams throughout the whole value chain from data scientists to SoC designers to work together. This design methodology must consider the needs of all stakeholders and provide both the required visibility to the design process and tools to handle the problems at different levels of the design chain. It must support selection of the optimal platform for the given algorithm as

well as the development of new electronics platforms for certain types of AI algorithms.

Algorithm selection and allocation problem

AI algorithms are usually developed in data centre environments, where the data scientists have practically no resource limitations. Processing latency is not critical because the focus is on the accuracy of the algorithm. Packaging algorithms into containers as in Figure 1 is simple and distinct.

Conversely in the edge device everything is limited and on the top of that the real-time requirements dictate the maximum latency of the computation. The ECUs in the vehicle platform network are usually optimised for certain types of tasks. They contain multiple processors and possibly dedicated accelerators for speeding up known algorithms like convolution or matrix multiplication. When functions of the algorithm are allocated to the existing platform network, some functions that logically belong together may not be executed in the same ECU but distributed across the platform network as in Figure 2.

The distributed implementation doesn't only distribute the functional workload to the different computing platform but also moves data communication between the functions from the local bus into the platform network. Multiple SW stacks running on different processing platforms and possibly different operating systems make the system latencies unpredictable and increase the risk of HW dependent software SW failures. Typical HW

failures are CPU overload, memory overload, bus congestion and data loss, which can cause SW malfunction or, in the worst case, even a system crash. The current trend in automotive industry, to move away from the device specific hardware architecture to a zonal architecture using High-Performance Computers (HPC), demands a very thorough corner case analysis with a multitude of scenarios.

In the zonal architecture the function allocation becomes a critical task. In the traditional design methodology, the function allocation is made by system architects based on their experience from previous projects, careful profiling and spreadsheet analysis. Yet, a Software Defined Vehicle (SDV) faces scenarios that nobody can imagine. Scenarios like "what happens in the emergency brake system, when the radar detects a pedestrian in front of the car, camera is blinded by the sun, some wheels are sliding on a slippery road and back seat entertainment system is downloading a video" are imaginative, but in real life possible – even probable.

Analysis of the different scenarios can lead to re-allocation of the functions to optimise the computing resource utilisation and platform network traffic, but consequently it can require re-architecting of the whole platform network including the accelerator architecture of the HPCs. Alternatively, the problematic algorithms can be optimised to better fit into the existing HW architecture. This process requires a seamless cooperation between data scientists, system architects,

SW developers and HW developers across the value chain, which includes the OEM and multiple tiers specialising in different technologies.

A new holistic design methodology must provide a common platform that all stakeholders can use and new methods to analyze the HW/SW performance metrics in a very early phase of the design process. The multi-tier value chain sets additional requirements to the design data handling like information sharing and hiding, requirements tracing across the value chain and seamless communication between the stakeholders.

Engineering software defined systems

The proposed design methodology follows the principles of Product Lifecycle Management (PLM) and uses elements from multiple system design methodologies. The key elements of the methodology are:

- Requirements management
- Model-Based Systems Engineering
- Domain specific electronics implementation flows
- Digitally threaded verification

Requirements management

High-quality requirements are key to a successful project. Requirements must be clear and verifiable in a formal way. A good practice is to parameterize as many requirements as possible with numerical or logical parameters. This enables automated tracing of requirements fulfilment in the design process. Requirements management tool must also enable itemisation of top-level requirements with more detailed and parameterizable sub-requirements.

Model-based systems engineering

Modelling and analysis of complex systems like SDV needs a modelling methodology that is generic enough to describe any kind of system or subsystem using the same design framework to minimise the communication problems between the different design domains. Model-based Systems Engineering (MBSE) methodologies are becoming more and more popular, replacing the textual specifications. Advanced MBSE methodologies like ARCADIA enable complexity management by using subsystem hierarchy, information sharing and hiding with black box and white box modelling, and requirements and parameter propagation throughout the model hierarchy. This is necessary in the multi-tier environment, where both OEM and subcontractors want to share only the relevant information with the project partners.




Figure 1. Multiple connected algorithms modelled as containers




Figure 2. Function allocation in a platform network




Figure 3. Model-Based electronics systems engineering process

ARCADIA methodology needs additional property-based methodology to fulfil the design needs of cyber-physical systems. Different implementation technologies of subsystems require implementation specific information that must be inserted into the system model, where it is needed. Property Modelling Methodology (PMM) has elements that are suitable for this purpose. The combined methodology is described in Figure 3. This methodology can be used to model all electronics system layers. Links to the domain specific implementation flows are not shown in the diagram.

Domain specific implementation flows

Hierarchical, subsystem-based implementation using different technologies is characteristic to the electronics system design. On the top of the implementation hierarchy is the network level that contains the cabling for data communication, power supply or signalling and the computing enclosures. The next level is the computing

enclosure (ECU or module) containing one or more printed circuit boards (PCB). Each PCB contains components that can be subsystems like SoC, FPGA or 3D integrated circuit. These can contain IP blocks or chiplets that are subsystems too.

The proposed engineering methodology supports multi-layer, multi-domain system modelling, where all components on each system layer are treated as components, although some of them can be decomposed further in separate subsystem projects and some of them are off-the-shelf parts. Therefore, each subsystem can be linked to a different implementation technology without causing any confusion on the upper subsystem layer.

The domain specific implementation flows are standard Electronics Design Automation (EDA) processes using conventional design tools. The input from the system modeling depends on the implementation flow and information available in the system model.






Figure 4. AI algorithm exploration and implementation flow

Digitally threaded verification

Developing complex systems like vehicles needs careful tracing of requirements, their allocation, implementation and verification. Conventional requirements tracing methodologies are not able to handle the complexity of the SW defined systems. Therefore, a digital threading methodology that links together the requirements, models, implementations, and verification from system analysis to system integration, is needed.

Optimising AI implementation

One special case of the software defined system engineering is the allocation of AI algorithms to the processing units. The system design process begins with requirements capture and analysis that can be done with a dedicated requirements management tool or using a spreadsheet. The software defined system engineering methodology enables a systematic analysis and optimisation of AI algorithm implementation. The process is described in the Figure 4.

The AI algorithms are developed in the neural network development environment like Tensorflow, PyTorch or Keras. The networks are optimised for accuracy and complexity and trained in this environment with the available data. Once the training is completed, the code is translated to C++ for implementation and analysis purposes. After the translation, the functions are profiled to estimate the workload on the processor.

In the next step the layered architecture is converted to a functional architecture of the system. In the case of multiple AI algorithms, all of them are included in the functional architecture. Following the ARCADIA methodology, the functions are allocated to logical components based on the profiling information and the logical components are allocated to the processing elements in the physical architecture shown in Figure 5. This initial allocation is based on the profiling information.

A performance simulation model is created based on this architecture model. The allocation to individual processor cores must be done in this phase. The ECU architecture can be taken into account in this phase. If the ECU contains several SoCs with multiple processors, the initial allocation to the individual

processors can be made here. An example simulation model is shown in Figure 6.

The performance simulation model can be parameterized to model different processor families, clock frequencies and bus widths. The communication channels are parameterizable for data size and trigger conditions.

The SW runnables can contain either a real snippet of the translated C code or a synthetic workload value. The synthetic workload approach enables very fast exploration of different accelerator configurations without changing the underlying model. The results are not as accurate as with the real source, but accurate enough for the initial analysis. The performance simulation models the execution of the SW tasks on the specified




Figure 5. Initial function allocation to processing units in multiple ECUs




Figure 6. Performance simulation model

HW architecture including the data communication delays, operating system effects, network utilisation and many other metrics. The simulation results can be visualised as HW and SW traces or statistics view.

Based on the performance simulation results, the function allocation or platform architecture can be changed to better utilise the resources available in the ECUs and platform network. Once the solution space is narrowed down to 2-3 potential implementation architectures, a virtual platform model can be created to run the real SW on the virtual HW model. The standard components like processors and memories are taken from the library and the custom accelerators are created by wrapping the functional C++ code into a SystemC wrapper. These models are not timing accurate, but they can be made bit-accurate by using bit-accurate C++ data types like `sc_fixed` or `sc_int` and updated with more accurate versions, when the timing accurate models are available.

Custom HW accelerator development needs some additional design steps. The C++ functional models must be optimised for High-Level Synthesis (HLS) and validated against the Python model.

HLS enables synthesis of C++ or SystemC code to VHDL or Verilog that is optimised for the target technology. Because the same source code is used for both HW and SW development, the functions can be freely allocated to SW tasks and HW accelerators and easily swapped between HW and SW even in the later phases of the design project. This enables the HW size and power consumption to be optimised by moving functions that are expensive to implement in HW to SW and taking HW friendly but computing intensive functions to the HW side.

The software defined system engineering methodology provides a comprehensive set of capabilities that enables implementation of any kind of AI functionality using existing ECU architectures or custom designs. The methodology is scalable and can be used for any design size and complexity. Digital twin and digitally threaded verification extend the methodology to a complete AI system design suite that supports multi-tier value chains.

INSIDE Members Focus

The micro-nano facility at the Fondazione Bruno Kessler

A sustainable model for innovation in the Chips Act era



Lorenza Ferrario



Since 1990, a semiconductor cleanroom located at the **Fondazione Bruno Kessler** (<https://www.fbk.eu>), in north Italy, has been supporting challenging research projects, SMEs and international agencies fabricating cutting edge silicon detectors and MEMS whose applications span medical to high energy physics, space, industry and the environment.

Over the years, huge investments have been optimised to constantly update the available technology and to keep, and improve, an excellence role at international level. And today, this cleanroom is part of the **FBK Micro-Nano Facility (MNF)**, including three other cleanrooms dedicated to silicon device fabrication, characterisation and dicing setups as well as a growing packaging and integration area, all suited to up to six-inch wafers. Courageous organisational choices made it possible to transform this reality from a “classic public research laboratory” to a pilot line that today enjoys a prestigious international standing and participates in strategic actions such as **IPCEI projects** and **Pilot Line Chips Act** side by side with key players in the European chips arena.

MNF: From R&D to pilot line

Since the very beginning of the MNF adventure, in 1990, all the building blocks of semiconductor device fabrication have been available in the first MNF cleanroom: thermal processing, ion implantation, lithography, dry and wet etching, deposition of thin film, online metrology. This has enabled researchers to play and train with this powerful “CMOS like” toys. And, in less than 2 years first CCD devices were successfully fabricated. The activity was organised according to classic R&D key performance indicators: publications, conferences, research projects; results possible thanks to the availability of some good optical sensors fabricated in the cleanroom.

Very soon, it appeared that running costs and investments associated with the maintenance of a cleanroom were high, not sustainable only with institutional funds, or budget coming from research projects, generally funding person months, travel expenses, and some consumables. Cleanrooms and their equipment must be functioning 24 hours

a day, all days of the year, regardless of their real use. Some materials used in the fabrication of chips are highly expensive. And so are process and technical gases needed to operate equipment. Equipment maintenance costs are high too. It is not easy to retain highly skilled staff, unless you can guarantee good salaries. Serious questions about the sustainability of the MNF first cleanroom were raised in those years.

Fortunately, in the late 1990s, something happened which completely re-addressed the Facility. We entered CERN tenders for the fabrication of large area radiation detectors, and double side radiation detectors (devices with aligned structures present on both sides of the wafer). Those devices were required for upgrades of experiments, like ALICE and ATLAS. Further opportunities to develop custom radiation sensors arrived thanks to projects with ESA/ASI and INFN which brought MNF devices mounted on the ISS. A huge challenge, since the requests were for several tens of devices, all of them fulfilling very narrow challenging specifications, in



some cases "large area detectors" which are extremely delicate and thus exceptionally challenging from the handling point of view. Last, but not least, all these projects had rigid deadlines, because of the nature of the experiments for which the devices were requested.

In order to meet all these demands, it was necessary to "forget behaving as pure researchers", working on "proof of concept" devices, more or less alone, and start familiarising with concepts like repeatability, statistical process control, process flow engineering and time control. It was also necessary to rethink professional roles by placing alongside the then central figures, researchers, with their know-how on devices and process flows, technologists and technicians, to be specialised on individual machines, not only for use but for maintenance, because it was necessary to guarantee high uptime and process control.

The aforementioned projects were successfully completed, and paved the way to complete and consolidate the "Facility" model, whose main features can be summarised as follows:

- Ability to fabricate small series of complete, controlled, reproducible devices. This for sure is what really allows MNF to support innovation, giving to companies the possibility to test markets with prototype series, or completely serve the market with niche products.
- Finding an alternative to "Moore's law", to limit the need for continuous equipment updates to fulfil the reduction of the technological nodes (today, commercial

processors and memories are in the range of some tens of nanometers). This has been possible since becoming a leader in devices which are appealing even with critical dimensions in the range 400-1000 nm. The CERN devices are a clear example: what is required is the quality, and the possibility to develop custom, and complex, design, not something that can be required for mass production foundries. Still today, we have lively activities in HEP and space projects requiring a CD in the order of 400 nm. Other examples of sensors and MEMS developed with this CD in MNF, and attracting collaborations and projects, are SiPM, Silicon photo-multiplier, gas sensors, flow sensors, custom photodiodes and phototransistors for industrial applications.

- ISO certification, quality based organisation. We happily realised that it is possible to manage under ISO 9001-2015 both services and R&D, that working according to ISO rules does not limit in any way the flexibility in the use of laboratories. On the contrary, it greatly helps in reducing the number of errors, for example, and so the time to complete fabrication, and the research in general. Thanks to ISO certification, first awarded to MNF in 2012, we widened opportunities to work with industry and with space agencies.
- Technology platforms suited to rapid customisation. While developing process flows to fabricate devices, MNF always try to define a general flow, as robust and short as possible, and suited to be adapted to different applications thanks

to the customisation of some parts of the fabrication steps. Two clear examples of this philosophy are silicon photomultiplier, SiPM, and MEMS microheaters. FBK SiPMs are now available basically for the whole light spectrum, from 100 nm (UV) to near infrared. FBK MEMS microheaters, based on a single design, are today used as components for chemoresistive gas sensors as well as flow sensors. The same flexibility is available on photodiodes, silicon drift detectors and some photonics components. In general, we dedicate great effort in designing customisable process flows, in order to significantly reduce the development time, but also boost the robustness of processes, and reliability of devices.

- Cost models. Besides scientific and technological skills, managing a complex facility like MNF requires knowledge and control of costs: cost of a device, of a lot, of an hour of tool use, of an hour of work of a researcher or technician. The control of this information allows the building of tariffs which guarantee to cover costs, and also, when collaborating with private companies, to correctly behave in order to comply with State Aid rules. In general, having a cost model makes it possible to correctly "price" not just the person months in EU projects, but also service activities. Furthermore, MNF is an Open Facility, and this means that external users, both researchers and companies, can access its capabilities. Cost models allow access to be transparently managed. Currently, MNF has two main types of costs: hourly / day rates for the use of back end facilities, and the "move" for pricing the front end, the cleanroom fabrication. A move is each process step done on each single wafer. So, a thermal oxidation done on 3 wafers is counted as 3 moves; an optical inspection done on 1 wafer is one move. A batch of 20 wafers requiring 100 fabrication steps corresponds to 2000 moves. What's the value of a move? Each year FBK updates the value of the move considering the cost of staff working in the cleanrooms, consumables, running costs directly connected to the cleanroom activity, utilities (electric energy, methane for heating, water), other direct components. Managing this information allows cost coverage to be controlled and this contributes highly to sustainability of the Facility.
- Moving from the "naked" sensor to the system. Especially in the last ten

years, the front end has been further enhanced, adding two other fabrication cleanrooms to the original one. But attention has been paid to the back end, in order to offer not just wafers and single devices, but packaged devices, and in some cases complete demo systems. As a result, the Facility today includes packaging and integration laboratories. The whole value chain is available at one side, from the wafer to the system ready for application use or test.

MNF technology roadmap

MNF is able to manufacture complete silicon sensors and MEMS enabling the development of applications in numerous areas: physics experiments, space, communications, medical,

bio, environment, agriculture, industry, renewables. At the Italian level, it is the only RTO capable of manufacturing complete silicon sensors in small series, with ISO certification, with TRL of up to 6. The chosen model is certainly a winner because it guarantees self-financing of over 75%, allowing the Facility to grow continuously over the last 25 years, and has laid the foundations for the "roadmap" that will accompany it for the next 10-15 years.

Thanks to its strong technology transfer orientation and in particular for the real capability to develop qualified small series of devices, 2019 MNF was admitted as a direct partner in the first IPCEI Microelectronics. The project enabled the expansion of manufacturing technology by adding slice thinning and temporary and final bonding processes for 3D integration developments. MNF took part in the formation of IPCEI Microelectronics ME/CT, the second IPCEI dedicated to microelectronics, in which it is, like other RTOs, an associate partner with its own proposal and budget for the construction of a new cleanroom entirely dedicated to Silicon Carbide and Germanium and thus to the development of a new generation of sensors based on these materials.

Principal technologies and capabilities

The main technologies, and devices, developed today in MNF are:

- Silicon radiation detectors, including SiPM, SDD, LGAD, pixel detectors, photodiodes.
- MEMS based sensors and actuators, including flow sensors, gas sensors, RF Switch, cMUT actuators, accelerometers
- Wearable sensors based on flexible electronics.
- photonics integrated circuits (PIC), including components in SiN, SiON, SOI, quantum photonics.
- CMOS technology, 3D vision for LiDAR
- Single photon sensors.
- Flex bonding.
- Integration of payloads for nanosatellites.

These devices and systems are used in innumerable applications, spanning medical, environmental, space, communications, AI, quantum sensing, quantum computing, and high energy physics experiments.

The Facility today incorporates by more than 2000 sqm of laboratories, including front end and back end capabilities. Here some technical details:

- Clean Room Detectors - complete silicon

sensors fabrication line. 700 m²; ISO Class 4/5 0,8 um CMOS pilot line: Ion Implantation, Oxidation, Diffusion, RIE, Deep RIE (silicon and oxide), Lithography (stepper 0.45 um, mask aligner, EBL), metal sputtering, SEM, inline metrology. Automatic handling for six-inch wafers.

- Clean Room MEMS – dedicated MEMS processing. 500 m² ISO Class 5/6, diffusion, lithography (mask aligner, DLW), wafer bonding, electroplating, Si bulk micromachining, metal evaporation, RIE, mechanical and optical profilometry, FIB-SEM nanolab for nanofabrication of QT and plasmonic. Automatic handling for six-inch wafers.
- Clean Room 3D IPCEI ME. 180 m² ISO Class 4/5 with wafer thinning and bonding technology for 3D integration wafer to wafer or chip to wafer, bonding sensors and sensor + readout electronics. Automatic handling for six-inch wafers
- Testing Area. 300 m² manual parametric testing, automatic parametric/functional testing, optical testing (spectral responsivity, quantum efficiency).
- Packaging and Integration Area. 100 m² clean room ISO Class 7 Microassembly station; screen printing, bonding (ball & wedge bonder), Shear-Pull Tester, reflow oven, CNC, pick and place: jet ball.
- In progress: cleanroom IPCEI ME/CT. 500 m² clean room ISO Class 4/5. SiC and Ge deposition, etching, thermal processes, 150 nm lithography, ion implantation for SiC substrates.



Image credit: Baroni

New Member Focus



David Baca



Jan Becker

Pioneering the future of mobility

How Apex.AI powers
real-time, safe and scalable vehicle
software

In the fast-evolving landscape of autonomous mobility and software-defined vehicles (SDVs), real-time responsiveness, safety and scalability are no longer technical luxuries, they're core requirements. At the heart of this transformation is Apex.AI, a technology leader that engineers deterministic, high-performance platforms built for the complexities of modern mobility. This new Inside Industry Association member explains its role in the cutting edge of edge computing, safety compliance and the software architectures that are driving the automotive industry into its next chapter.

Edge intelligence: computing at the speed of real-world decision making

From LIDAR arrays to AI-driven path planning, autonomous systems must process massive volumes of data with split-second accuracy. Apex.AI's answer is edge computing with real-time performance baked in. The company's flagship product, Apex.OS, is designed for flexible deployment across heterogeneous hardware and OS configurations while supporting industry protocols like DDS, CAN and SOME/IP. Coupled with Apex.Ida, a zero-copy, shared-memory middleware, the platform enables high-throughput communication with minimal CPU overhead. This is vital for handling the deluge of sensor data in safety-critical environments. "Determinism is not a feature, it's a requirement" is the Apex.AI conviction, whose Grace Executor ensures predictable, repeatable execution for every process running in a vehicle, whether it's handling diagnostics or dynamic route planning.

Safety-certified AI for autonomous systems

With AI taking centre stage in mobility systems, compliance with safety standards like ISO 26262 and SOTIF is mission-critical. Apex.AI's platforms are not just compliant but are built from the ground up for functional safety. Deterministic execution, system health monitoring and replay mechanisms enable rapid anomaly detection. Apex.OS and Apex.Grace are developed to meet ASIL D safety levels, providing OEMs the peace of mind needed for real-world deployment of intelligent, autonomous capabilities.

Middleware that moves the industry forward

What separates Apex.Ida from conventional middleware is its ability to maintain multi-gigabit throughput with minimal latency while integrating across legacy and modern systems. The platform is protocol-agnostic, bridging DDS,

CAN, MQTT, SOME/IP and more under a unified API. This drastically simplifies integration and streamlines performance optimisation across ECUs.

Virtual validation: from sim to street

Testing in virtual environments is now a critical step in developing safe ADAS and ADS features. Apex.OS supports MCAP and TECMP formats, integrating with leading simulators like Carla and dSpace to ensure deterministic behaviour under test conditions. After all, simulation isn't just about cost savings – it's about confidence, repeatability and finding issues before you hit the road. In simple terms, the Apex.OS helps customers close the gap between simulation and real-world deployment.

From ROS 2 to road-ready production

One of Apex.AI's most developer-friendly features is its ROS 2 API compatibility. While ROS 2 remains the industry standard for prototyping, it lacks the safety and real-time guarantees needed for production. Apex.Grace bridges this gap by retaining the ROS 2 interface but delivering an under-the-hood transformation that meets rigorous production standards. This enables OEMs to transition prototypes into products without reinventing their software stack.

Enabling sustainability through smart architectures

Europe's green transition is not just about EVs; it's about efficiency across the entire vehicle platform. Apex.AI supports ECU consolidation and mixed-criticality workloads, reducing hardware redundancy and optimising energy use. This aligns with sustainable manufacturing goals and reduces long-term software maintenance burdens.

Powering the software-defined vehicle

Apex.AI supports the future of software-defined vehicles (SDVs) through a robust software platform tailored for real-time, safe and

scalable vehicle software development. These capabilities position Apex.AI as a foundational platform for SDV architectures, supporting the evolution of vehicles into secure, reliable, and software-centric systems. SDVs represent the next generation of mobility: connected, intelligent, upgradable. Apex.AI is building the foundation with a platform that supports:

- **Real-Time and Safety-Critical Execution:**

Apex.Grace provides deterministic, real-time capable execution with the Apex.Grace Executor and execution monitoring tools. It supports precise control over threads, core affinities and scheduling, essential for reliable SDV operation.

- **Scalable Communication Middleware:**

Apex.Ida offers high-performance, low-latency communication with zero-copy IPC, support for various automotive protocols (CAN, DDS, SOME/IP, MQTT), and a flexible connector architecture for seamless integration across legacy and modern systems.

- **Software Portability and ROS**

Compatibility: Apex.Grace maintains API compatibility with ROS 2, allowing reuse of existing nodes and tools while providing enhanced safety and deterministic guarantees. This accelerates transition from R&D to production.

- **System Health, Diagnostics and**

Lifecycle Management: Apex.Grace includes a Process Manager, System Manager and diagnostics stack (UDS, DoIP, DTC handling) for orchestrated process control and fault management across complex systems.

The road ahead for mobility is defined not just by hardware, but by the software that orchestrates it all. Apex.AI is a conductor of this orchestration, leading that transformation of how vehicles are built, validated and evolved.

Contact: info@apex.ai

New Member Focus



Driving impact beyond research

Innovation Dis.Co and the power
of strategic exploitation



Vicky Chatzidogiannaki

In the dynamic world of EU-funded research and innovation, turning technical excellence into lasting societal and economic impact remains a central challenge. At Innovation Dis.Co, we embrace that challenge as our core mission. As a boutique consultancy specialized in Communication, Dissemination, and Exploitation (CDE) for Horizon Europe and Chips JU projects, we go beyond compliance, delivering tailored strategies that drive real adoption, visibility, and value.

With a diverse background spanning social sciences, design, data analytics and technology, our team brings a holistic view to each project. From Smart Mobility to Electronic Components and Systems (ECS), we help translate complex research into narratives, tools and stakeholder journeys that make sense, and make impact.

This article explores our approach to project exploitation, highlighting our unique Exploitation Radar methodology and digital tools that ensure research doesn't end with a deliverable, but leads to real-world uptake.

Our role in EU projects: not just communication but strategic positioning

While many organisations treat communication and dissemination as parallel tracks to technical

work, at Innovation Dis.Co we integrate CDE activities directly into the heart of the project's value proposition. Our services span:

- Strategic communication planning and branding.
- Sentiment Analysis monitoring.
- Research impact analysis – dissemination performance tracking.
- Stakeholder mapping and engagement
- Market and competitive landscape intelligence.
- Exploitation planning and Key Exploitable Result (KER) management.
- Event management (workshops, review meetings, clustering activities).

To embed CDE activities at the very heart of every project, we designed the Innovation Dis.Co CDE Platform — an integrated suite of three complementary pillars:




Figure 1: Innovation Dis.Co Exploitation RoadMap Radar




Figure 2: Innovation Dis.Co Sentiment Analysis Platform

- Exploitation Radar (Exploitation)
- Research Impact Monitor (Dissemination)
- Sentiment Analysis Monitor (Communication)

In the sections that follow, we unpack our hands-on approach for each module.

Exploitation is not an afterthought: it's a roadmap

In many EU projects, exploitation is still treated as a box-ticking exercise, addressed late in the project or reduced to generic templates and vague intentions. This traditional approach often fails to create real momentum for adoption, ownership or sustainability. At Innovation Dis.Co, we take a different view: exploitation is not a final step, it's a continuous, strategic process that should begin on day one. By embedding exploitation planning early and revisiting it throughout the project lifecycle, we help ensure that results evolve in parallel with stakeholder needs, market opportunities, and policy relevance, leading to outcomes that are not just documented, but truly impactful.

Our methodology helps consortia:

- Identify and refine Key Exploitable Results (KERs).
- Align technical outcomes with stakeholder and market needs.
- Assess the competitive landscape.
- Plan realistic paths to uptake, whether commercial, policy-oriented or societal.
- Distinguish between KERs at different Technology Readiness Levels (TRLs), enabling the prioritization of those closer to market maturity and focusing exploitation efforts on results that can deliver immediate value.

To support this approach, we developed a unique monitoring and facilitation tool: the *Exploitation Radar*.


The exploitation radar: a digital framework for real-time uptake readiness

The Exploitation Radar is a dynamic platform that allows project coordinators, partners and stakeholders to monitor the maturity, relevance and uptake potential of their results in real time.

Key features include:

- A visual dashboard that maps KERs by domain, application readiness and exploitation route.
- Real-time filtering by sector (mobility, energy, health, manufacturing, etc.)
- Integration with exploitation questionnaire(s) responses to build dynamic profiles.
- Automatic clustering suggestions for internal or external synergies.
- A public-facing version of the Platform, where any external stakeholder can explore project results, via an interactive map, filter by country or organisation, and access detailed profiles of each partner's Key Exploitable Results — covering TRLs, core technologies and statistical insights. This version is designed to promote visibility and foster new collaboration opportunities across industry, academia and policymakers, all while protecting sensitive partner data through tiered disclosure controls.

The Exploitation Radar is not only a tool, it's a methodology. It underpins our workshops, guides our partner interactions and supports a culture of early ownership and strategic planning.





Recent case highlights: from insight to impact

We have successfully applied the Exploitation Radar across multiple EU-funded projects, including:

- **PowerizeD:** Facilitating the identification and prioritization of Key Exploitable Results in power electronics, enabling clearer market orientation and enhanced collaboration among industrial and academia stakeholders.
- **EcoMobility:** Supporting the market uptake of cooperative, connected, and automated mobility (CCAM) technologies through real-time exploitation planning and stakeholder mapping.
- **ShapeFuture:** Helping ECS partners identify pathways for technology transfer, standardisation and industrial uptake.

In each case, our contribution ensures that exploitation is not just documented, but acted upon — with tangible plans for sustainability, ownership and further funding.

Recognized and evolving through feedback

Throughout our involvement in EU-funded projects, our exploitation methodology, has received consistently positive feedback from project reviewers and officers. Their input has helped us refine our tools and adapt our approach to better serve the evolving needs of research and innovation consortia. We view this recognition not as an end point, but as motivation to continue developing practical, transparent and stakeholder-oriented exploitation strategies across all domains.

Communication expertise: more than messaging

At Innovation Dis.Co, our communication work goes far beyond basic outreach. We act as a full-service communication agency within EU-funded projects, crafting compelling narratives, visuals, and campaigns that elevate the project's identity and ensure meaningful engagement across all audiences. We develop impactful branding strategies tailored to each project's objectives and ecosystem, ensuring visibility and clarity from day one.

Our approach is data-driven, supported by our custom-built Sentiment Analysis Platform, which monitors stakeholder perception across digital channels in real time. This insight allows us to adjust messaging, optimize timing, and focus communication efforts where they matter most. Combined with creative design and audience mapping, we

deliver dynamic and effective communication plans that support both dissemination and exploitation goals.

Research impact monitor: monitoring scientific impact

Monitoring scientific impact in a Horizon Europe project is key to meeting EU dissemination obligations and demonstrating the value of research outputs. It helps track visibility through publications, citations, and online engagement; informs stakeholders; supports communication and exploitation activities; and improves outreach strategies. Ultimately, it ensures the project's results are used, recognized and impactful beyond the consortium.

To that end, the Research Impact Monitor supports:

- **Citation analysis:** Tracking and analyzing total citations, downloads and readership.
- **Online engagement monitoring:** Capturing research mentions across social media, news websites, policy documents and blogs.
- **Visualization:** Displaying the dissemination footprint of project research in real time.

Conclusion: scaling impact, one result at a time

Innovation Dis.Co exists to empower research teams to think beyond the project. By combining strategic insight, human-centric design and digital tooling, we offer a unique value proposition in the EU innovation ecosystem. We welcome collaboration with project coordinators, industrial players, policy stakeholders and other CDE specialists to scale the impact of research and innovation across Europe.

If you're ready to turn your results into real-world outcomes, let's connect.

www.innovation-disco.com

New Member Focus

Crossing Europe's valley of death

Why semiconductor innovation needs a product path



Ozgur Ozkurt

Europe has made major strides in semiconductor policy. Through the Chips Act, we've seen critical investments in design capability, pilot lines and R&D infrastructure. But for all this momentum, a core structural weakness remains: Europe lacks robust mechanisms to carry semiconductor ventures across the *Valley of Death* - the dangerous gap between a successful research demonstrator and a commercially viable product.

At Quintauris, we operate right at this inflection point, developing platforms that span real-time to high-performance compute, system integration and the enabling software that supports hardware-software ecosystems. We are not alone. Across Europe, a new generation of scale-ups is emerging - technically ambitious, strategically aligned, and ready to deliver industrial value. A key enabler for this is the adoption of open standards, like RISC-V, which are critical for Europe's long-term technological ownership and prevent vendor lock-in. At Quintauris, our platforms are built exclusively on this open standard, reflecting our commitment to a sovereign and collaborative ecosystem. Still, many ventures face structural hurdles when transitioning from prototype to product.

Several recurring patterns explain why this gap persists:

- The Pilot-to-Product Gap. Many publicly funded projects stop short of deployment. Even successful demonstrators at high TRLs often lack a clear path to broad usability or scalable integration. Without a plan for production alignment or reuse across multiple applications, their impact remains isolated.
- The Software Enablement Gap. To be truly usable, new hardware must be complemented by robust software tooling, open SDKs and clear integration pathways. Shared device farms and cloud access are critical enablers, especially for smaller players who rely on early adoption to innovate.
- The Customer-Integration Gap. Ventures that engage closely with end-users - across automotive, industrial or AI domains - can adapt early and validate faster. Without this proximity, promising technologies struggle to find their market fit.

The challenge is not lack of innovation - it is lack of continuity.



The journey from early R&D to a validated, reusable platform requires long-term support, not just seed funding. It requires a mindset shift: to treat ecosystem-building and platform scalability as core to Europe's competitiveness.

To complement the foundations already laid by the Chips Act, Europe's next step must be to build the bridge across this valley. We propose a focus on:

1. **Dedicated Productization Tracks**

We need funding instruments, perhaps as dedicated Chips JU Focus Topic calls,

for scale-ups working at TRL7 and above to help them mature demonstrators into market-ready products.

2. **Ecosystem Enablement Funding**

Support must extend beyond the chip itself to include the software communities, developer tools and open platforms that drive adoption and create a thriving ecosystem.

3. **Customer-Led Innovation Programmes**

Prioritise and support ventures that are co-developing solutions directly with lead customers, ensuring that innovation is shaped by real-world deployment realities from day one.

4. **Agile Funding for Growth**

Europe must diversify its investment methodology. Alongside traditional consortium grants, we need agile, venture-style funding tracks for scale-ups needing to grow. This "shark tank" approach - with lean applications and a focus on commercial scaling - would empower companies without paralysing them with processes designed for R&D consortia.

The ambition is shared across the Union. The technical capability exists within our institutions and companies. What's needed now is a policy framework that bridges invention and execution - one that empowers the teams building not just technologies, but the platforms upon which Europe's future industries will be built.

New Member Focus

Cubit Innovation Labs



Lorenzo Monti

"At Cubit, we believe innovation must be accessible, applicable, and human-centered. By combining cutting-edge research with real industrial needs, we aim to build not just technologies, but solutions that matter."

- Lucio Russo, CEO Cubit Innovation Labs

Cubit is a public-private consortium headquartered at the Technology Park of Navacchio, near Pisa, Italy. Cubit was founded in 2007 by the Department of Information Engineering of the University of Pisa, the Technology Park of Navacchio, and several highly innovative Italian companies with the objective of bridging technological research and industrial expertise, promoting innovation and technology transfer in the B2B domain.

Cubit is dedicated to driving innovation in the industrial sector through research, development, and high-tech engineering services. By bridging the gap between academia and business, Cubit plays a key role in the Italian and European innovation ecosystem.

With its multidisciplinary team - composed of more than 35 employees - and strong collaboration with the University of Pisa, Cubit provides customized R&D and product engineering services to industrial entities from various market sectors - from automotive and aerospace to industry and sports applications - to fulfill their product and process innovation needs.

Cubit's operational model supports the entire innovation pipeline, covering:

1. Feasibility & design: identifying needs, scouting technologies, and defining specifications
2. Prototyping & simulation: developing the first version of the prototype or the Digital Twin to quickly validate its functionality, performance, and market potential.
3. Validation & testing: EMC, environmental, mechanical and RF compliance
4. Industrialization support: integrating the product into the company's strategy, also providing the necessary consulting, expertise, and training.

This full-cycle approach enables both startups and large enterprises to bring complex technologies to market faster, with reduced risk and greater agility.

Areas of expertise

Cubit operates through two core divisions:

■ IoT division

Cubit designs and develops customized IoT devices and embedded platforms, enabling real-time data acquisition, wireless communication, and intelligent

control. Its expertise spans hardware and firmware design, cloud integration, cybersecurity, and edge AI. Applications include industrial automation, smart energy, building monitoring, wearable tech, people and asset tracking, and environmental sensing. IoT division supports clients from concept to certification, ensuring a fast, robust, and cost-effective path to market.

■ Fluid dynamics division

This division specializes in advanced computational fluid dynamics (CFD), optimization, and experimental validation. Working in close collaboration with professors and researchers from the Department of Aerospace Engineering at the University of Pisa, the team uses a 9216-core computing cluster and proprietary simulation tools to optimize flow dynamics in sectors such as marine engineering, automotive, and aerospace.

In 2013, Cubit established a dedicated branch in Maranello to serve a major automotive client on-site, offering full-time CFD expertise embedded within the customer's design process.

Cubit also provides rapid prototyping through in-house 3D printing and has access to certified laboratories, including a semi-anechoic chamber for EMC measurements.

Applied research and strategic vision

At the heart of Cubit's mission is the belief that research and innovation must be accessible, applied, and future-oriented. Cubit acts as a catalyst for industrial innovation by translating cutting-edge research into practical engineering outcomes. Thanks to its consorcial structure and strong academic roots, Cubit combines methodological rigor with operational flexibility, allowing it to move quickly between concept, design, and testing.

Cubit's strategy is based on three key pillars:

- Research-driven development, fostering continuous exchange between academic knowledge and real-world engineering
- Cross-sector innovation, transferring solutions across domains such as health, mobility, energy, and marine tech
- Collaborative networks, working closely with European partners, competence centers, and consortia to co-develop technologies and scale impact.

This model enables Cubit not only to offer advanced engineering services, but also to contribute to the development of a European innovation ecosystem that values interdisciplinarity, openness, and sustainability.

Collaboration in European R&D projects

As a member of INSIDE Industry Association and an active player in European research frameworks, Cubit contributes to several ongoing projects that reflect its commitment to applied innovation:

- **DUCA:** Trust management and cybersecurity in intelligent digital networks, with a focus on explainability and zero-trust paradigms.
- **EU-TRAINS:** AI-based wearable textiles for real-time health and performance monitoring, combining sensor innovation, machine learning, and digital health.
- **H2TRAIN:** development of next-generation biosensors for digital health and lifestyle monitoring, combining advanced sensing materials with low-power embedded systems and AI for real-time physiological and environmental data acquisition.

Success story: NUVAP - advanced indoor air quality monitoring

Among Cubit's most successful long-term collaborations is NUVAP, a technology



startup that developed an innovative solution for indoor air quality monitoring.

Launched with support from Italy's Smart&Start Italia program, funded by the Ministry of Economic Development (MISE) and Invitalia, the project received over €1 million to bring to market a device capable of detecting and analyzing multiple pollutants in indoor environments. The system helps assess the healthiness of living and working spaces, enabling both individuals and organisations to take preventive or corrective actions.

Since 2014, Cubit has been a strategic partner throughout the entire innovation journey, contributing to:

- The initial concept and feasibility
- Design and prototyping of custom sensors (including a proprietary radon sensor)
- Integration of multiple sensing modules into a compact, connected device
- Certification, compliance testing, and final product engineering

NUVAP exemplifies how deep collaboration between startups and applied research centers can turn visionary ideas into certified, market-ready technologies that support both sustainability and public health.



Figure 1: Cubit's FPGA board

What's next

As innovation cycles accelerate and the demand for intelligent, sustainable technologies continues to rise, Cubit is looking ahead with a strong commitment to both technological depth and societal relevance.

One of the key directions in Cubit's R&D roadmap is the development of custom FPGA-based platforms. With the increasing demand for high-performance, low-latency computing at the edge, FPGAs have become essential in domains such as industrial automation, automotive, medical devices, and signal processing.

Cubit's in-house board, CORE-SF2 & EVB, is the result of years of experience in embedded systems and real-time control. Designed to support rapid prototyping and scalable deployment, it provides engineers and researchers with a versatile and robust toolset for creating complex architectures with a high degree of customization and reliability.

From AI accelerators to secure edge computing, Cubit's FPGA platforms are developed to meet the specific challenges of today's embedded applications while remaining adaptable to tomorrow's needs.



Figure 2: Musai prototype

Parallel to that, thanks to the European projects H2TRAIN and EU-TRAINS, Cubit is specializing in smart textiles and wearable technologies, expanding its expertise in sensor integration, embedded AI, and low-power electronics. These efforts are not only aimed at health monitoring or sport performance, but also at emerging domains such as ambient sensing, workplace safety, and assistive technologies—where comfort, discretion, and data precision are essential. Another key area for growth is the blue economy, where Cubit applies its fluid dynamics, embedded systems, and sustainability know-how to new solutions for marine monitoring and ocean protection.

This is the case of the project MUSAI – Modular Underwater System with Artificial Intelligence - a modular underwater system equipped with artificial intelligence and advanced sensor technologies, developed to optimize ocean monitoring. Designed to operate even at significant depths, MUSAI is capable of identifying, classifying, and localizing marine pollution, processing information directly on-edge and transmitting data in real time to a support platform. Thanks to its modularity, in addition to waste detection, the system is also suitable for other research applications, such as monitoring the health of marine flora. MUSAI stands as a concrete technological solution, designed to support research institutions, environmental operators, and companies committed to protecting marine ecosystems.

Across all these initiatives, Cubit remains focused on its core mission: to connect cutting-edge research with real-world needs, enabling industries to innovate with confidence, purpose, and measurable impact.

ENSURE Workshop

October 6-9th
2025
Patras Greece

Chairs

Markus Tauber
Research Studios Austria

Paolo Azzoni
INSIDE Industry Association

Gianluigi Ferrari
University of Parma

Amin Anjomshoaa
*Vienna University of
Economics and Business*

TPC Members

Dimitrios Serpanos
Janis Judvaitis
George Nikolakopoulos
Stefan Van Baelen
Andrea Pinna
Perez Juan Martin
Paolo Azzoni
Gianluigi Ferrari
Markus Tauber
Amin Anjomshoaa
Jerker Delsing
Thomas Felberbauer
Christian Jandl

Important dates

Submission: 04/08/2025
Acceptance: 11/08/2025
Camera ready: 18/08/2025
Workshop: 06/10/2025

at IEEE International Smart Cities Conference – ISC2 2025

The ENSURE (Embedding iNtelligence in Smart mUnicipalities – challenges and oppoRtunitiEs) workshop will be a one day workshop at the ISC2 conference. It brings together researchers and engineers from academia and industry to identify and solve critical issues in the rapid technological and regulatory transformations shaping the future of urban environments. The convergence of technologies such as Artificial Intelligence (AI) at the edge, Cyber-Physical Systems of Systems (CPSoS), and the IoT-edge-cloud continuum is rapidly reshaping how urban services are designed and delivered. However, these advancements do not come without risks: cybersecurity, privacy, compliance with evolving regulatory frameworks (such as the EU AI Act, Data Governance Act, and Cyber Resilience Act), and ethical concerns must be addressed from the start. This workshop stands out because it integrates both technological and regulatory perspectives. Participants will discuss not only the latest technological innovations but also how to embed intelligence in a way that respects citizen rights, enhances trust, and complies with regulations: ENSURE brings a holistic, system-of-systems view to urban intelligence, considering societal, ethical, and regulatory impacts.

Topics of interest include, but are not limited to:

- Embedding intelligence in smart municipalities
- Edge AI
- IoT-edge-cloud continuum for urban environments
- Smart Cyber Physical Systems of Systems
- Security, Privacy and Compliance with regulatory frameworks
- Distributed and pervasive systems for smart cities
- Trusted data spaces and federated data infrastructures for smart cities
- Embedded and cyberphysical systems for smart cities, regions and municipalities
- Industrial systems for smart municipality infrastructures
- Applications and services in smart municipalities
- ML and AI in smart environments
- Emerging network technologies for smart municipalities
- Society 5.0
- Legal and regulatory aspects of smart cities, regions and municipalities

Submissions of Papers

All accepted papers will be hosted on IEEE Xplore as peer-reviewed archival. Submissions are expected to not have been published or submitted elsewhere.

Full technical papers must be formatted using the IEEE 2-column format and not exceed 6 pages for full paper submissions or not exceed 4 pages for short paper submissions. Short papers can be positioning papers and work in progress reports which are also very much welcome.



For
submissions
follow this link
or scan the
QR code

Research Project Highlight

AIMS5.0

Advancing AI in manufacturing for
a sustainable Industry 5.0



Thomas Gutt

Through AI innovation and strategic collaboration, AIMS5.0 is laying the foundation for Industry 5.0, an era where sustainability, digital leadership and human-centred innovation converge to redefine manufacturing across Europe and beyond.

In today's fast-changing industrial landscape, Europe is leading a transformative movement that combines digital sovereignty, environmental responsibility and technological innovation. At the heart of this transformation lies AIMS5.0, a €70 million collaborative Innovation Action that brings together 53 leading academic and industry partners to redefine the future of manufacturing.

From Industry 4.0 to Industry 5.0

While Industry 4.0 focused on automation and data exchange, Industry 5.0 takes a decisive step further, prioritizing human-centricity, resilience and sustainability.

Through the integration of cutting-edge technologies, such as Artificial Intelligence (AI), Machine Learning (ML), the Internet of Things (IoT) and Semantic Web Ontologies, AIMS5.0 aims to create smart factories that are not only more productive but also climate-friendly and socially responsible.

The project's vision is clear: to strengthen European digital sovereignty and boost competitiveness by deploying AI-enabled hardware and software systems across the entire industrial value chain.

This shift marks a significant milestone in the continent's strategy to retain and grow its production capabilities within Europe, ensuring shorter supply chains, enhanced resilience and greater sustainability.

Transforming semiconductor manufacturing

One of the most striking applications of this initiative can be seen in semiconductor manufacturing, a sector critical to the global economy and digital infrastructure. Through AI-powered defect image classification, manufacturers can now identify imperfections in real time, dramatically increasing production efficiency, reducing waste and ensuring consistently high-quality output.



By improving accuracy and reducing waste, these AI-driven advancements streamline operations while significantly lowering environmental impact. AIMS5.0 is developing innovative chip technologies and tools that span the entire value chain, enabling faster time-to-market and broader acceptance of new, versatile AI-driven technologies.

How AI is transforming semiconductor manufacturing

AIMS5.0 has unveiled a brand new explainer video demonstrating how AI is revolutionising semiconductor manufacturing. Viewers

can explore real-world examples of AI applications in defect detection, resource optimisation and sustainable production.

Watch the video on the AIMS5.0 YouTube channel to see these groundbreaking technologies in action. Discover how AI-powered solutions are making chip production more efficient, sustainable and precise.

This short explainer video from AIMS5.0 showcases:

- Real-time defect detection through image classification
- Smarter resource use and reduced environmental impact
- Cutting-edge examples of Industry 5.0 technology in manufacturing

A human-centric, climate-friendly future

What sets AIMS5.0 apart is its commitment to human-centered digitalisation. The project is not just about machines getting smarter, it's about creating workplaces that are safer, more inclusive and better aligned with human values. By placing people and the planet at the core of innovation, AIMS5.0 is helping shape an industrial future that balances economic growth with environmental stewardship.

AIMS5.0: pioneering Europe's sustainable manufacturing

European initiative is more than just a technological leap, it's a movement toward a greener, smarter and more resilient manufacturing ecosystem. To learn more about the project and its partners, visit:

- AIMS5.0 Official Site
- Infineon Technologies
- KAI Competence Center

Together, through AI innovation and collaborative excellence, AIMS5.0 is building the foundation of Industry 5.0, where sustainability, technology and humanity converge.

CEI-Sphere

Strengthening Europe's digital fabric through market insights and stakeholder dialogue





Veronica Vuotto

Nine months into its journey, the CEI-Sphere project continues to lay critical groundwork for Europe's Cloud-Edge-IoT (CEI) ecosystem. Launched in October 2024 under the Horizon Europe programme, CEI-Sphere plays an important role in coordinating efforts to build an easy, interoperable, and resilient CEI infrastructure across the continent. While the Large-Scale Pilots (LSPs) become more mature and ready after their kick-off last January, the project has already made substantial progress in equipping stakeholders with the knowledge and tools they need to navigate a fast-changing digital landscape.

From its beginning, CEI-Sphere has positioned itself not only as a facilitator of technological integration but also as a knowledge hub. The last few weeks have for instance seen the release of the first Market Briefs, a fundamental element of the project's ambition to find a connection between technological innovation and market realities.

These briefs, now available on the project website, are the result of comprehensive market analyses aimed at mapping out current trends, challenges, and opportunities across key industries. Organised by key sectors, they provide specific insights for stakeholders operating in distinct yet interconnected domains.

Each brief explains complex dynamics into clear, simple content, offering readers a better understanding of how Cloud-Edge-IoT technologies play an important role in critical sectors such as mobility, energy, and manufacturing.

Take, for instance, the brief on CEI in the Mobility Sector, which explores how global supply chain disruptions and the growing need for visibility are fuelling demand for real-time, interconnected solutions. In the Energy Sector, CEI-Sphere's analysis addresses the compounded challenges of geopolitical volatility, rising costs, and climate imperatives, an area where CEI technologies could make a profound impact by enhancing digital resilience and enabling decarbonisation. Meanwhile, the brief on Manufacturing illustrates how real-time data flows and AI-driven automation



are becoming essential for maintaining competitiveness in an increasingly turbulent economic environment.

Beyond individual sectors, the briefs also tackle broader themes such as emerging and disruptive technologies, shedding light on how tech convergence is driving innovation in CEI, and the intricate nature of the CEI ecosystem itself, where collaboration across hardware, software, data, and security layers is essential for scalable and secure deployment.

To further open up these insights to a broader audience and spark dialogue, CEI-Sphere has launched a dedicated series of "CEI Market Talks". These webinars bring together experts, innovators, and policymakers to explore the findings of each Market Brief in greater depth, offering a platform for discussion, exchange, and collective sense-making.

- The first session, held on 30 May, focused on Mobility, a sector undergoing rapid digital and structural change. As supply chains face mounting pressure from global disruptions, CEI technologies are enabling greater visibility, efficiency, and resilience, particularly through real-time data flows and predictive logistics solutions.

- This was followed by the second Market Talk, which took place on 12 June, and turned the spotlight on the Energy sector. Here, the conversation explored how geopolitical instability, high energy prices, and the push for decarbonisation are driving the need for smarter, more agile energy systems. CEI solutions are helping accelerate this transition, supporting the integration of renewables, enabling demand-response models, and boosting system-wide sustainability.
- The series will conclude on 16 July with a session dedicated to Manufacturing, a domain where digital transformation is no longer optional, but essential for competitiveness and survival. European manufacturers are increasingly turning to CEI-powered innovations to overcome inflation, global competition, and operational volatility.


Looking ahead, these activities support the foundation of CEI-Sphere's wider engagement strategy. With the two large-scale pilots progressing their work, these early milestones ensure that the CEI conversation is not just technical, but also market-aware, and deeply grounded in the real-world needs of European industries.

The months ahead promise to be equally rich in insights and action. CEI-Sphere is indeed set to play an important role in shaping Europe's digital and innovative future.


CEI-Sphere will be joining the first edition of INSIDE Connect, taking place on 3–4 September 2025. We're looking forward to meeting fellow innovators, showcasing our work, and contributing to the conversations about Europe's digital future. See you in Málaga!

Follow us on Social Media!

Let's get connected for our latest news, events and updates!



 /inside-association



INSIDE
Industry Association

Online version is available at Inside-association.eu

Publisher

INSIDE Industry Association
High Tech Campus 69-3
5656 AG Eindhoven, The Netherlands

Design and Creative lay-out


Studio Kraft – Veldhoven, the Netherlands

Acknowledgements

With thanks to the interviewees, project participants, INSIDE Industry Association office, the INSIDE Industry Association Presidium and other INSIDE Industry Association-involved persons for any assistance and material provided in the production of this issue of the INSIDE Magazine.

Contributions



The INSIDE Industry Association office is interested in receiving news or events in the field of Intelligent Digital Systems. Please submit your information to info@Inside-association.eu



INSIDE
Industry Association

© 2025 INSIDE Industry Association

Permission to reproduce individual articles from INSIDE Magazine for non- commercial purposes is granted, provided that INSIDE Magazine is credited as the source. Opinions expressed in the INSIDE Magazine do not necessarily reflect those of the organisation.



INSIDE
Industry Association

[INSIDE-association.eu](https://inside-association.eu)